vce trainer

AWS Certified Solutions Architect - Professional.302q

Architect-Prof

AWS-Architect

AWS Certified Solutions Architect – Professional

Score: 800/1000

Version:

Licensed to Berend Post berend.post@it-university.nl

Time Limit: 120 Minutes

Licensed to Berend Post berend.post@it-university.nl

Exam A (302 questions)

Question 1

By default, what is the maximum number of Cache Nodes you can run in Amazon ElastiCache?

- ☐ 20
- ☐ 50
- ☐ 100
- ☐ 200

Question 2

In Amazon ElastiCache, which of the following statements is correct?

- ☐ When you launch an ElastiCache cluster into an Amazon VPC private subnet, every cache node is assigned a public IP address within that subnet.
- ☐ You cannot use ElastiCache in a VPC that is configured for dedicated instance tenancy.
- ☐ If your AWS account supports only the EC2-VPC platform, ElastiCache will never launch your cluster in a VPC.
- ☐ ElastiCache is not fully integrated with Amazon Virtual Private Cloud (VPC).

Question 3

Which of the following is the Amazon Resource Name (ARN) condition operator that can be used within an Identity and Access Management (IAM) policy to check the case-insensitive matching of the ARN?

- ☐ ArnCheck
- ☐ ArnMatch
- ☐ ArnCase
- ☐ ArnLike

Question 4

A user is creating a PIOPS volume. What is the maximum ratio the user should configure between PIOPS and the volume size?

- ☐ 5
- ☐ 10
- ☐ 20
- ☐ 30

Licensed to Berend Post berend.post@it-university.nl

Question 5

Which statement is NOT true about accessing remote AWS region in the US by your AWS Direct Connect which is located in the US?

- ☐ AWS Direct Connect locations in the United States can access public resources in any US region.
- ☐ You can use a single AWS Direct Connect connection to build multi-region services.
- ☐ Any data transfer out of a remote region is billed at the location of your AWS Direct Connect data transfer rate.
- ☐ To connect to a VPC in a remote region, you can use a virtual private network (VPN) connection over your public virtual interface.

Question 6

What feature of the load balancing service attempts to force subsequent connections to a service to be redirected to the same node as long as it is online?

- ☐ Node balance
- ☐ Session retention
- ☐ Session multiplexing
- ☐ Session persistence

Question 7

An organization is planning to host a Wordpress blog as well a joomla CMS on a single instance launched with VPC. The organization wants to have separate domains for each application and assign them using Route 53. The organization may have about ten instances each with two applications as mentioned above. While launching the instance, the organization configured two separate network interfaces (primary + ENI) and wanted to have two elastic IPs for that instance. It was suggested to use a public IP from AWS instead of an elastic IP as the number of elastic IPs is restricted.

What action will you recommend to the organization?

- ☐ I agree with the suggestion but will prefer that the organization should use separate subnets with each ENI for different public IPs.
- ☐ I do not agree as it is required to have only an elastic IP since an instance has more than one ENI and AWS does not assign a public IP to an instance with multiple ENIs.
- ☐ I do not agree as AWS VPC does not attach a public IP to an ENI; so the user has to use only an elastic IP only.
- ☐ I agree with the suggestion and it is recommended to use a public IP from AWS since the organization is going to use DNS with Route 53.

Licensed to Berend Post berend.post@it-university.nl

Question 8

You have just added a new instance to your Auto Scaling group, which receives ELB health checks. An ELB heath check says the new instance's state is out of Service.

What does Auto Scaling do in this particular scenario?

- ☐ It replaces the instance with a healthy one
- ☐ It stops the instance
- ☐ It marks an instance as unhealthy
- ☐ It terminates the instance

Question 9

An organization has hosted an application on the EC2 instances. There will be multiple users connecting to the instance for setup and configuration of application. The organization is planning to implement certain security best practices.

Which of the below mentioned pointers will not help the organization achieve better security arrangement?

- ☐ Allow only IAM users to connect with the EC2 instances with their own secret access key.
- ☐ Create a procedure to revoke the access rights of the individual user when they are not required to connect to EC2 instance anymore for the purpose of application configuration.
- ☐ Apply the latest patch of OS and always keep it updated.
- ☐ Disable the password based login for all the users. All the users should use their own keys to connect with the instance securely.

Question 10

Amazon Elastic File System (EFS) provides information about the space used for an object by using the space _ used attribute of the Network File System Version 4.1 (NFSv4.1). The attribute includes the object's current metered data size and not the metadata size. Which of the following utilities will you use to measure the amount of disk that is used of a file?

- ☐ blkid utility
- ☐ du utility
- ☐ sfdisk utility
- ☐ pydf utility

Licensed to Berend Post berend.post@it-university.nl

Question 11

Out of the striping options available for the EBS volumes, which one has the following disadvantage: 'Doubles the amount of I/O required from the instance to EBS compared to RAID 0, because you're mirroring all writes to a pair of volumes, limiting how much you can stripe.'?

- ☐ Raid 1
- ☐ Raid 0
- ☐ RAID 1+0 (RAID 10)
- ☐ Raid 2

Question 12

A company runs a three-tier application in AWS. Users report that the application performance can vary greatly depending on the time of day and functionality being accessed.

The application includes the following components:
Eight t2.large front-end web servers that serve static content and proxy dynamic content from the application tier.
Four t2.large application servers.
One db.m4.large Amazon RDS MySQL Multi-AZ DB instance.

Operations has determined that the web and application tiers are network constrained.

Which of the following should cost effective improve application performance? (Choose two.)

- ☐ Replace web and app tiers with t2.xlarge instances
- ☐ Use AWS Auto Scaling and m4.large instances for the web and application tiers
- ☐ Convert the MySQL RDS instance to a self-managed MySQL cluster on Amazon EC2
- ☐ Create an Amazon CloudFront distribution to cache content
- ☐ Increase the size of the Amazon RDS instance to db.m4.xlarge

Licensed to Berend Post berend.post@it-university.nl

Question 13

You deployed your company website using Elastic Beanstalk and you enabled log file rotation to S3. An Elastic Map Reduce job is periodically analyzing the logs on S3 to build a usage dashboard that you share with your CIO.
You recently improved overall performance of the website using Cloud Front for dynamic content delivery and your website as the origin.
After this architectural change, the usage dashboard shows that the traffic on your website dropped by an order of magnitude.

How do you fix your usage dashboard?

- ☐ Enable Cloud Front to deliver access logs to S3 and use them as input of the Elastic Map Reduce job.
- ☐ Turn on Cloud Trail and use trail log tiles on S3 as input of the Elastic Map Reduce job
- ☐ Change your log collection process to use Cloud Watch ELB metrics as input of the Elastic Map Reduce job
- ☐ Use Elastic Beanstalk "Rebuild Environment" option to update log delivery to the Elastic Map Reduce job.
- ☐ Use Elastic Beanstalk "Restart App server(s)" option to update log delivery to the Elastic Map Reduce job.

Licensed to Berend Post berend.post@it-university.nl

Question 14

A corporate web application is deployed within an Amazon Virtual Private Cloud (VPC) and is connected to the corporate data center via an IPSec VPN. The application must authenticate against the on-premises LDAP server. After authentication, each logged-in user can only access an Amazon Simple Storage Space (S3) keyspace specific to that user.

Which two approaches can satisfy these objectives? (Choose 2)

- ☐ Develop an identity broker that authenticates against IAM security Token service to assume a IAM role in order to get temporary AWS security credentials The application calls the identity broker to get AWS temporary security credentials with access to the appropriate S3 bucket.
- ☐ The application authenticates against LDAP and retrieves the name of an IAM role associated with the user. The application then calls the IAM Security Token Service to assume that IAM role. The application can use the temporary credentials to access the appropriate S3 bucket.
- ☐ Develop an identity broker that authenticates against LDAP and then calls IAM Security Token Service to get IAM federated user credentials. The application calls the identity broker to get IAM federated user credentials with access to the appropriate S3 bucket.
- ☐ The application authenticates against LDAP the application then calls the AWS identity and Access Management (IAM) Security service to log in to IAM using the LDAP credentials the application can use the IAM temporary credentials to access the appropriate S3 bucket.
- ☐ The application authenticates against IAM Security Token Service using the LDAP credentials the application uses those temporary AWS security credentials to access the appropriate S3 bucket.

Licensed to Berend Post berend.post@it-university.nl

Question 15

You are designing a multi-platform web application for AWS The application will run on EC2 instances and will be accessed from PCs. Tablets and smart phones Supported accessing platforms are Windows, MacOS, IOS and Android Separate sticky session and SSL certificate setups are required for different platform types.

Which of the following describes the most cost effective and performance efficient architecture setup?

- ☐ Setup a hybrid architecture to handle session state and SSL certificates on-prem and separate EC2 Instance groups running web applications for different platform types running in a VPC.
- ☐ Set up one ELB for all platforms to distribute load among multiple instance under it Each EC2 instance implements ail functionality for a particular platform.
- ☐ Set up two ELBs The first ELB handles SSL certificates for all platforms and the second ELB handles session stickiness for all platforms for each ELB run separate EC2 instance groups to handle the web application for each platform.
- ☐ Assign multiple ELBS to an EC2 instance or group of EC2 instances running the common components of the web application, one ELB for each platform type Session stickiness and SSL termination are done at the ELBs.

Question 16

What bandwidths do AWS Direct Connect currently support?

- ☐ 10Mbps and 100Mbps
- ☐ 10Gbps and 100Gbps
- ☐ 100Mbps and 1Gbps
- ☐ 1Gbps and 10 Gbps

Licensed to Berend Post berend.post@it-university.nl

Question 17

A bank is designing an online customer service portal where customers can chat with customer service agents. The portal is required to maintain a 15-minute RPO or RTO in case of a regional disaster. Banking regulations require that all customer service chat transcripts must be preserved on durable storage for at least 7 years, chat conversations must be encrypted in-flight, and transcripts must be encrypted at rest. The Data Lost Prevention team requires that data at rest must be encrypted using a key that the team controls, rotates, and revokes.

Which design meets these requirements?

- ☐ The chat application logs each chat message into Amazon CloudWatch Logs. A scheduled AWS Lambda function invokes a CloudWatch Logs. CreateExportTask every 5 minutes to export chat transcripts to Amazon S3. The S3 bucket is configured for cross-region replication to the backup region. Separate AWS KMS keys are specified for the CloudWatch Logs group and the S3 bucket.
- ☐ The chat application logs each chat message into two different Amazon CloudWatch Logs groups in two different regions, with the same AWS KMS key applied. Both CloudWatch Logs groups are configured to export logs into an Amazon Glacier vault with a 7-year vault lock policy with a KMS key specified.
- ☐ The chat application logs each chat message into Amazon CloudWatch Logs. A subscription filter on the CloudWatch Logs group feeds into an Amazon Kinesis Data Firehose which streams the chat messages into an Amazon S3 bucket in the backup region. Separate AWS KMS keys are specified for the CloudWatch Logs group and the Kinesis Data Firehose.
- ☐ The chat application logs each chat message into Amazon CloudWatch Logs. The CloudWatch Logs group is configured to export logs into an Amazon Glacier vault with a 7-year vault lock policy. Glacier cross-region replication mirrors chat archives to the backup region. Separate AWS KMS keys are specified for the CloudWatch Logs group and the Amazon Glacier vault.

Question 18

A user has created a VPC with public and private subnets using the VPC wizard. Which of the below mentioned statements is true in this scenario?

- ☐ The user has to manually create a NAT instance
- ☐ The Amazon VPC will automatically create a NAT instance with the micro size only
- ☐ VPC updates the main route table used with the private subnet, and creates a custom route table with a public subnet
- ☐ VPC updates the main route table used with a public subnet, and creates a custom route table with a private subnet

Licensed to Berend Post berend.post@it-university.nl

Question 19

A company has an application written using an in-house software framework. The framework installation takes 30 minutes and is performed with a user data script. Company Developers deploy changes to the application frequently. The framework installation is becoming a bottleneck in this process.

Which of the following would speed up this process?

- ☐ Create a pipeline to build a custom AMI with the framework installed and use this AMI as a baseline for application deployments.
- ☐ Employ a user data script to install the framework but compress the installation files to make them smaller.
- ☐ Create a pipeline to parallelize the installation tasks and call this pipeline from a user data script.
- ☐ Configure an AWS OpsWorks cookbook that installs the framework instead of employing user data. Use this cookbook as a base for all deployments.

Question 20

The Principal element of an IAM policy refers to the specific entity that should be allowed or denied permission, whereas the translates to everyone except the specified entity.

- ☐ NotPrincipal
- ☐ Vendor
- ☐ Principal
- ☐ Action

Question 21

An organization is planning to setup a management network on the AWS VPC. The organization is trying to secure the webserver on a single VPC instance such that it allows the internet traffic as well as the back-end management traffic. The organization wants to make so that the back end management network interface can receive the SSH traffic only from a selected IP range, while the internet facing webserver will have an IP address which can receive traffic from all the internet IPs.

How can the organization achieve this by running web server on a single instance?

- ☐ It is not possible to have two IP addresses for a single instance.
- ☐ The organization should create two network interfaces with the same subnet and security group to assign separate IPs to each network interface.
- ☐ The organization should create two network interfaces with separate subnets so one instance can have two subnets and the respective security groups for controlled access.
- ☐ The organization should launch an instance with two separate subnets using the same network interface which allows to have a separate CIDR as well as security groups.

Licensed to Berend Post berend.post@it-university.nl

Question 22

The user has provisioned the PIOPS volume with an EBS optimized instance.
Generally speaking, in which I/O chunk should the bandwidth experienced by the user be measured by AWS?

- ☐ 128 KB
- ☐ 256 KB
- ☐ 64 KB
- ☐ 32 KB

Question 23

An organization has created 5 IAM users. The organization wants to give them the same login ID but different passwords. How can the organization achieve this?

- ☐ The organization should create each user in a separate region so that they have their own URL to login
- ☐ The organization should create a separate login ID but give the IAM users the same alias so that each one can login with their alias
- ☐ It is not possible to have the same login ID for multiple IAM users of the same account
- ☐ The organization should create various groups and add each user with the same login ID to different groups. The user can login with their own group ID

Question 24

Can a user configure a custom health check with Auto Scaling?

- ☐ Yes, but the configured data will not be saved to Auto Scaling.
- ☐ No, only an ELB health check can be configured with Auto Scaling.
- ☐ Yes
- ☐ No

Licensed to Berend Post berend.post@it-university.nl

Question 25

A company runs a public-facing application that uses a Java-based web service via a RESTful API. It is hosted on Apache Tomcat on a single server in a data center that runs consistently at 30% CPU utilization. Use of the API is expected to increase by 10 times with a new product launch. The business wants to migrate the application to AWS with no disruption, and needs it to scale to meet demand.

The company has already decided to use Amazon Route 53 and CNAME records to redirect traffic. How can these requirements be met with the LEAST amount of effort?

- ☐ Use AWS Elastic Beanstalk to deploy the Java web service and enable Auto Scaling. Then switch the application to use the new web service.
- ☐ Lift and shift the Apache server to the cloud using AWS SMS. Then switch the application to direct web service traffic to the new instance.
- ☐ Create a Docker image and migrate the image to Amazon ECS. Then change the application code to direct web service queries to the ECS container.
- ☐ Modify the application to call the web service via Amazon API Gateway. Then create a new AWS Lambda Java function to run the Java web service code. After testing, change API Gateway to use the Lambda function.

Question 26

A web-startup runs its very successful social news application on Amazon EC2 with an Elastic Load Balancer, an Auto-Scaling group of Java/Tomcat application-servers, and DynamoDB as data store. The main web-application best runs on m2 x large instances since it is highly memory- bound Each new deployment requires semi-automated creation and testing of a new AMI for the application servers which takes quite a while ana is therefore only done once per week.
Recently, a new chat feature has been implemented in nodejs and wails to be integrated in the architecture. First tests show that the new component is CPU bound Because the company has some experience with using Chef, they decided to streamline the deployment process and use AWS Ops Works as an application life cycle tool to simplify management of the application and reduce the deployment cycles.

What configuration in AWS Ops Works is necessary to integrate the new chat module in the most cost-efficient and flexible way?

- ☐ Create one AWS OpsWorks stack, create one AWS Ops Works layer, create one custom recipe
- ☐ Create one AWS OpsWorks stack create two AWS Ops Works layers, create one custom recipe
- ☐ Create two AWS OpsWorks stacks create two AWS Ops Works layers, create one custom recipe
- ☐ Create two AWS OpsWorks stacks create two AWS Ops Works layers, create two custom recipe

Licensed to Berend Post berend.post@it-university.nl

Question 27

When does an AWS Data Pipeline terminate the AWS Data Pipeline-managed compute resources?

- ☐ AWS Data Pipeline terminates AWS Data Pipeline-managed compute resources every 2 hours.
- ☐ When the final activity that uses the resources is running
- ☐ AWS Data Pipeline terminates AWS Data Pipeline-managed compute resources every 12 hours.
- ☐ When the final activity that uses the resources has completed successfully or failed

Question 28

ABC has created a multi-tenant Learning Management System (LMS). The application is hosted for five different tenants (clients) in the VPCs of the respective AWS accounts of the tenant. ABC wants to setup a centralized server which can connect with the LMS of each tenant upgrade if required. ABC also wants to ensure that one tenant VPC should not be able to connect to the other tenant VPC for security reasons.

How can ABC setup this scenario?

- ☐ ABC has to setup one centralized VPC which will peer in to all the other VPCs of the tenants.
- ☐ ABC should setup VPC peering with all the VPCs peering each other but block the IPs from CIDR of the tenant VPCs to deny them.
- ☐ ABC should setup all the VPCs with the same CIDR but have a centralized VPC. This way only the centralized VPC can talk to the other VPCs using VPC peering.
- ☐ ABC should setup all the VPCs meshed together with VPC peering for all VPCs.

Question 29

You require the ability to analyze a large amount of data, which is stored on Amazon S3 using Amazon Elastic Map Reduce. You are using the cc2 8xlarge instance type, whose CPUs are mostly idle during processing. Which of the below would be the most cost efficient way to reduce the runtime of the job?

- ☐ Create more, smaller flies on Amazon S3.
- ☐ Add additional cc2 8xlarge instances by introducing a task group.
- ☐ Use smaller instances that have higher aggregate I/O performance.
- ☐ Create fewer, larger files on Amazon S3.

Licensed to Berend Post berend.post@it-university.nl

Question 30

True or False: Amazon ElastiCache supports the Redis key-value store.

- ☐ True, ElastiCache supports the Redis key-value store, but with limited functionalities.
- ☐ False, ElastiCache does not support the Redis key-value store.
- ☐ True, ElastiCache supports the Redis key-value store.
- ☐ False, ElastiCache supports the Redis key-value store only if you are in a VPC environment.

Question 31

An Auto Scaling group is running at the desired capacity of 5 instances and receives a trigger from the Cloudwatch Alarm to increase the capacity by 1. The cool down period is 5 minutes. Cloudwatch sends another trigger after 2 minutes to decrease the desired capacity by 1.

What will be the count of instances at the end of 4 minutes?

- ☐ 4
- ☐ 5
- ☐ 6
- ☐ 7

Licensed to Berend Post berend.post@it-university.nl

Question 32

A company is moving a business-critical, multi-tier application to AWS. The architecture consists of a desktop client application and server infrastructure. The server infrastructure resides in an on-premises data center that frequently fails to maintain the application uptime SLA of 99.95%. A Solutions Architect must re-architect the application to ensure that it can meet or exceed the SLA.

The application contains a PostgreSQL database running on a single virtual machine. The business logic and presentation layers are load balanced between multiple virtual machines. Remote users complain about slow load times while using this latency-sensitive application.

Which of the following will meet the availability requirements with little change to the application while improving user experience and minimizing costs?

☐ Migrate the database to a PostgreSQL database in Amazon EC2. Host the application and presentation layers in automatically scaled Amazon ECS containers behind an Application Load Balancer. Allocate an Amazon WorkSpaces WorkSpace for each end user to improve the user experience.

☐ Migrate the database to an Amazon RDS Aurora PostgreSQL configuration. Host the application and presentation layers in an Auto Scaling configuration on Amazon EC2 instances behind an Application Load Balancer. Use Amazon AppStream 2.0 to improve the user experience.

☐ Migrate the database to an Amazon RDS PostgreSQL Multi-AZ configuration. Host the application and presentation layers in automatically scaled AWS Fargate containers behind a Network Load Balancer. Use Amazon ElastiCache to improve the user experience.

☐ Migrate the database to an Amazon Redshift cluster with at least two nodes. Combine and host the application and presentation layers in automatically scaled Amazon ECS containers behind an Application Load Balancer. Use Amazon CloudFront to improve the user experience.

Licensed to Berend Post berend.post@it-university.nl

Question 33

A Company has a security event whereby an Amazon S3 bucket with sensitive information was made public. Company policy is to never have public S3 objects, and the Compliance team must be informed immediately when any public objects are identified.

How can the presence of a public S3 object be detected, set to trigger alarm notifications, and automatically remediated in the future? (Choose two.)

- ☐ Turn on object-level logging for Amazon S3. Turn on Amazon S3 event notifications to notify by using an Amazon SNS topic when a PutObject API call is made with a public-read permission.
- ☐ Configure an Amazon CloudWatch Events rule that invokes an AWS Lambda function to secure the S3 bucket.
- ☐ Use the S3 bucket permissions for AWS Trusted Advisor and configure a CloudWatch event to notify by using Amazon SNS.
- ☐ Turn on object-level logging for Amazon S3. Configure a CloudWatch event to notify by using an SNS topic when a PutObject API call with public-read permission is detected in the AWS CloudTrail logs.
- ☐ Schedule a recursive Lambda function to regularly change all object permissions inside the S3 bucket.

Question 34

AWS Direct Connect itself has NO specific resources for you to control access to. Therefore, there are no AWS Direct Connect Amazon Resource Names (ARNs) for you to use in an Identity and Access Management (IAM) policy.

With that in mind, how is it possible to write a policy to control access to AWS Direct Connect actions?

- ☐ You can leave the resource name field blank.
- ☐ You can choose the name of the AWS Direct Connection as the resource.
- ☐ You can use an asterisk (*) as the resource.
- ☐ You can create a name for the resource.

Question 35

What is the maximum length for an instance profile name in AWS IAM?

- ☐ 512 characters
- ☐ 128 characters
- ☐ 1024 characters
- ☐ 64 characters

Licensed to Berend Post berend.post@it-university.nl

Question 36

Which statement is NOT true about a stack which has been created in a Virtual Private Cloud (VPC) in AWS OpsWorks?

- ☐ Subnets whose instances cannot communicate with the Internet are referred to as public subnets.
- ☐ Subnets whose instances can communicate only with other instances in the VPC and cannot communicate directly with the Internet are referred to as private subnets.
- ☐ All instances in the stack should have access to any package repositories that your operating system depends on, such as the Amazon Linux or Ubuntu Linux repositories.
- ☐ Your app and custom cookbook repositories should be accessible for all instances in the stack.

Question 37

An organization is planning to create a secure scalable application with AWS VPC and ELB. The organization has two instances already running and each instance has an ENI attached to it in addition to a primary network interface. The primary network interface and additional ENI both have an elastic IP attached to it.
If those instances are registered with ELB and the organization wants ELB to send data to a particular EIP of the instance, how can they achieve this?

- ☐ The organization should ensure that the IP which is required to receive the ELB traffic is attached to a primary network interface.
- ☐ It is not possible to attach an instance with two ENIs with ELB as it will give an IP conflict error.
- ☐ The organization should ensure that the IP which is required to receive the ELB traffic is attached to an additional ENI.
- ☐ It is not possible to send data to a particular IP as ELB will send to any one EIP.

Question 38

What is the default maximum number of VPCs allowed per region?

- ☐ 5
- ☐ 10
- ☐ 100
- ☐ 15

Licensed to Berend Post berend.post@it-university.nl

Question 39

Which of the following statements is NOT correct when working with your AWS Direct Connect connection after it is set up completely?

☐ You can manage your AWS Direct Connect connections and view the connection details.
☐ You can delete a connection as long as there are no virtual interfaces attached to it.
☐ You cannot view the current connection ID and verify if it matches the connection ID on the Letter of Authorization (LOA).
☐ You can accept a host connection by purchasing a hosted connection from the partner (APN).

Question 40

Which of the following rules must be added to a mount target security group to access Amazon Elastic File System (EFS) from an on-premises server?

☐ Configure an NFS proxy between Amazon EFS and the on-premises server to route traffic.
☐ Set up a Point-To-Point Tunneling Protocol Server (PPTP) to allow secure connection.
☐ Permit secure traffic to the Kerberos port 88 from the on-premises server.
☐ Allow inbound traffic to the Network File System (NFS) port (2049) from the on-premises server.

Licensed to Berend Post berend.post@it-university.nl

Question 41

You are the new IT architect in a company that operates a mobile sleep tracking application.
When activated at night, the mobile app is sending collected data points of 1 kilobyte every 5 minutes to your backend.
The backend takes care of authenticating the user and writing the data points into an Amazon DynamoDB table.
Every morning, you scan the table to extract and aggregate last night's data on a per user basis, and store the results in Amazon S3. Users are notified via Amazon SNS mobile push notifications that new data is available, which is parsed and visualized by the mobile app.
Currently you have around 100k users who are mostly based out of North America.
You have been tasked to optimize the architecture of the backend system to lower cost.

What would you recommend? (Choose 2)

- ☐ Have the mobile app access Amazon DynamoDB directly Instead of JSON files stored on Amazon S3.
- ☐ Write data directly into an Amazon Redshift cluster replacing both Amazon DynamoDB and Amazon S3.
- ☐ Introduce an Amazon SQS queue to buffer writes to the Amazon DynamoDB table and reduce provisioned write throughput.
- ☐ Introduce Amazon Elasticache to cache reads from the Amazon DynamoDB table and reduce provisioned read throughput.
- ☐ Create a new Amazon DynamoDB table each day and drop the one for the previous day after its data is on Amazon S3.

Question 42

In DynamoDB, a projection is_____.

- ☐ systematic transformation of the latitudes and longitudes of the locations inside your table
- ☐ importing data from your file to a table
- ☐ exporting data from a table to your file
- ☐ the set of attributes that is copied from a table into a secondary index

Licensed to Berend Post berend.post@it-university.nl

Question 43

Your system recently experienced down time during the troubleshooting process. You found that a new administrator mistakenly terminated several production EC2 instances.

Which of the following strategies will help prevent a similar situation in the future?

The administrator still must be able to:
launch, start stop, and terminate development resources.
launch and start production instances.

- ☐ Create an IAM user, which is not allowed to terminate instances by leveraging production EC2 termination protection.
- ☐ Leverage resource based tagging, along with an IAM user which can prevent specific users from terminating production, EC2 resources.
- ☐ Leverage EC2 termination protection and multi-factor authentication, which together require users to authenticate before terminating EC2 instances
- ☐ Create an IAM user and apply an IAM role which prevents users from terminating production EC2 instances.

Question 44

You are designing an intrusion detection prevention (IDS/IPS) solution for a customer web application in a single VPC. You are considering the options for implementing IOS IPS protection for traffic coming from the Internet.

Which of the following options would you consider? (Choose 2 answers)

- ☐ Implement IDS/IPS agents on each Instance running in VPC
- ☐ Configure an instance in each subnet to switch its network interface card to promiscuous mode and analyze network traffic.
- ☐ Implement Elastic Load Balancing with SSL listeners in front of the web applications
- ☐ Implement a reverse proxy layer in front of web servers and configure IDS/IPS agents on each reverse proxy server.

Question 45

Will you be able to access EC2 snapshots using the regular Amazon S3 APIs?

- ☐ Yes, you will be able to access using S3 APIs if you have chosen the snapshot to be stored in S3.
- ☐ No, snapshots are only available through the Amazon EBS APIs.
- ☐ Yes, you will be able to access them using S3 APIs as all snapshots are stored in S3.
- ☐ No, snapshots are only available through the Amazon EC2 APIs.

Licensed to Berend Post berend.post@it-university.nl

Question 46

An AWS customer is deploying an application mat is composed of an AutoScaling group of EC2 Instances.
The customers security policy requires that every outbound connection from these instances to any other service within the customers Virtual Private Cloud must be authenticated using a unique x 509 certificate that contains the specific instance-id.
In addition, an x 509 certificates must Designed by the customer's Key management service in order to be trusted for authentication.

Which of the following configurations will support these requirements?

- ☐ Configure an IAM Role that grants access to an Amazon S3 object containing a signed certificate and configure the Auto Scaling group to launch instances with this role. Have the instances bootstrap get the certificate from Amazon S3 upon first boot.
- ☐ Embed a certificate into the Amazon Machine Image that is used by the Auto Scaling group. Have the launched instances generate a certificate signature request with the instance's assigned instance-id to the key management service for signature.
- ☐ Configure the Auto Scaling group to send an SNS notification of the launch of a new instance to the trusted key management service. Have the Key management service generate a signed certificate and send it directly to the newly launched instance.
- ☐ Configure the launched instances to generate a new certificate upon first boot. Have the Key management service poll the Auto Scaling group for associated instances and send new instances a certificate signature (hat contains the specific instance-id.

Question 47

A company wants to replace its call system with a solution built using AWS managed services. The company call center would like the solution to receive calls, create contact flows, and scale to handle growth projections. The call center would also like the solution to use deep learning capabilities to recognize the intent of the callers and handle basic tasks, reducing the need to speak an agent. The solution should also be able to query business applications and provide relevant information back to calls as requested.

Which services should the Solution Architect use to build this solution? (Choose three.)

- ☐ Amazon Rekognition to identity who is calling.
- ☐ Amazon Connect to create a cloud-based contact center.
- ☐ Amazon Alexa for Business to build conversational interface.
- ☐ AWS Lambda to integrate with internal systems.
- ☐ Amazon Lex to recognize the intent of the caller.
- ☐ Amazon SQS to add incoming callers to a queue.

Licensed to Berend Post berend.post@it-university.nl

Question 48

In Amazon IAM, what is the maximum length for a role name?

- ☐ 128 characters
- ☐ 512 characters
- ☐ 64 characters
- ☐ 256 characters

Question 49

A company is using AWS for production and development workloads. Each business unit has its own AWS account for production, and a separate AWS account to develop and deploy its applications. The Information Security department has introduced new security policies that limit access for terminating certain Amazon ECs instances in all accounts to a small group of individuals from the Security team.

How can the Solutions Architect meet these requirements?

- ☐ Create a new IAM policy that allows access to those EC2 instances only for the Security team. Apply this policy to the AWS Organizations master account.
- ☐ Create a new tag-based IAM policy that allows access to these EC2 instances only for the Security team. Tag the instances appropriately, and apply this policy in each account.
- ☐ Create an organizational unit under AWS Organizations. Move all the accounts into this organizational unit and use SCP to apply a whitelist policy to allow access to these EC2 instances for the Security team only.
- ☐ Set up SAML federation for all accounts in AWS. Configure SAML so that it checks for the service API call before authenticating the user. Block SAML from authenticating API calls if anyone other than the Security team accesses these instances.

Question 50

Who is responsible for modifying the routing tables and networking ACLs in a VPC to ensure that a DB instance is reachable from other instances in the VPC?

- ☐ AWS administrators
- ☐ The owner of the AWS account
- ☐ Amazon
- ☐ The DB engine vendor

Licensed to Berend Post berend.post@it-university.nl

Question 51

A user is planning to use EBS for his DB requirement. The user already has an EC2 instance running in the VPC private subnet.
How can the user attach the EBS volume to a running instance?

- ☐ The user can create EBS in the same zone as the subnet of instance and attach that EBS to instance.
- ☐ It is not possible to attach an EBS to an instance running in VPC until the instance is stopped.
- ☐ The user can specify the same subnet while creating EBS and then attach it to a running instance.
- ☐ The user must create EBS within the same VPC and then attach it to a running instance.

Question 52

In Amazon ElastiCache, the failure of a single cache node can have an impact on the availability of your application and the load on your back-end database while ElastiCache provisions a replacement for the failed cache node and it get repopulated.

Which of the following is a solution to reduce this potential availability impact?

- ☐ Spread your memory and compute capacity over fewer number of cache nodes, each with smaller capacity.
- ☐ Spread your memory and compute capacity over a larger number of cache nodes, each with smaller capacity.
- ☐ Include fewer number of high capacity nodes.
- ☐ Include a larger number of cache nodes, each with high capacity.

Question 53

Which of the following are characteristics of Amazon VPC subnets? (Choose 2)

- ☐ Each subnet spans at least 2 Availability Zones to provide a high-availability environment.
- ☐ Each subnet maps to a single Availability Zone.
- ☐ CIDR block mask of /25 is the smallest range supported.
- ☐ By default, all subnets can route between each other, whether they are private or public.
- ☐ Instances in a private subnet can communicate with the Internet only if they have an Elastic IP.

Licensed to Berend Post berend.post@it-university.nl

Question 54

By default, temporary security credentials for an IAM user are valid for a maximum of 12 hours, but you can request a duration as long as _____ hours.

- ☐ 24
- ☐ 36
- ☐ 10
- ☐ 48

Question 55

A company has detected to move some workloads onto AWS to create a grid environment to run market analytics. The grid will consist of many similar instances, spun-up by a job-scheduling function. Each time a large analytics workload is completed, a new VPC is deployed along with job scheduler and grid nodes. Multiple grids could be running in parallel.

Key requirements are:

Grid instances must communicate with Amazon S3 retrieve data to be processed.
Grid instances must communicate with Amazon DynamoDB to track intermediate data,
The job scheduler need only to communicate with the Amazon EC2 API to start new grid nodes.

A key requirement is that the environment has no access to the internet, either directly or via the on-premises proxy. However, the application needs to be able to seamlessly communicate to Amazon S3, Amazon DynamoDB, and Amazon EC2 API, without the need for reconfiguration for each new deployment.

Which of the following should the Solutions Architect do to achieve this target architecture? (Choose three.)

- ☐ Enable VPC endpoints for Amazon S3 and DynamoDB.
- ☐ Disable Private DNS Name Support.
- ☐ Configure the application on the grid instances to use the private DNS name of the Amazon S3 endpoint.
- ☐ Populate the on-premises DNS server with the private IP addresses of the EC2 endpoint.
- ☐ Enable an interface VPC endpoint for EC2.
- ☐ Configure Amazon S3 endpoint policy to permit access only from the grid nodes.

Licensed to Berend Post berend.post@it-university.nl

Question 56

A company has developed a web application that runs on Amazon EC2 instances in one AWS Region. The company has taken on new business in other countries and must deploy its application into other to meet low-latency requirements for its users. The regions can be segregated, and an application running in one region does not need to communicate with instances in other regions.

How should the company's Solutions Architect automate the deployment of the application so that it can be MOST efficiently deployed into multiple regions?

- ☐ Write a bash script that uses the AWS CLI to query the current state in one region and output a JSON representation. Pass the JSON representation to the AWS CLI, specifying the --region parameter to deploy the application to other regions.
- ☐ Write a bash script that uses the AWS CLI to query the current state in one region and output an AWS CloudFormation template. Create a CloudFormation stack from the template by using the AWS CLI, specifying the --region parameter to deploy the application to other regions.
- ☐ Write a CloudFormation template describing the application's infrastructure in the resources section. Create a CloudFormation stack from the template by using the AWS CLI, specify multiple regions using the --regions parameter to deploy the application.
- ☐ Write a CloudFormation template describing the application's infrastructure in the Resources section. Use a CloudFormation stack set from an administrator account to launch stack instances that deploy the application to other regions.

Question 57

MapMySite is setting up a web application in the AWS VPC. The organization has decided to use an AWS RDS instead of using its own DB instance for HA and DR requirements.
The organization also wants to secure RDS access.

How should the web application be setup with RDS?

- ☐ Create a VPC with one public and one private subnet. Launch an application instance in the public subnet while RDS is launched in the private subnet.
- ☐ Setup a public and two private subnets in different AZs within a VPC and create a subnet group. Launch RDS with that subnet group.
- ☐ Create a network interface and attach two subnets to it. Attach that network interface with RDS while launching a DB instance.
- ☐ Create two separate VPCs and launch a Web app in one VPC and RDS in a separate VPC and connect them with VPC peering.

Licensed to Berend Post berend.post@it-university.nl

Question 58

AWS _____supports_____ environments as one of the AWS resource types.

- ☐ Elastic Beanstalk; Elastic Beanstalk application
- ☐ CloudFormation; Elastic Beanstalk application
- ☐ Elastic Beanstalk ; CloudFormation application
- ☐ CloudFormation; CloudFormation application

Question 59

You are designing a data leak prevention solution for your VPC environment. You want your VPC Instances to be able to access software depots and distributions on the Internet for product updates. The depots and distributions are accessible via third party CDNs by their URLs.
You want to explicitly deny any other outbound connections from your VPC instances to hosts on the internet.

Which of the following options would you consider?

- ☐ Configure a web proxy server in your VPC and enforce URL-based rules for outbound access Remove default routes.
- ☐ Implement security groups and configure outbound rules to only permit traffic to software depots.
- ☐ Move all your instances into private VPC subnets remove default routes from all routing tables and add specific routes to the software depots and distributions only.
- ☐ Implement network access control lists to all specific destinations, with an Implicit deny all rule.

Question 60

What is a possible reason you would need to edit claims issued in a SAML token?

- ☐ The NameIdentifier claim cannot be the same as the username stored in AD.
- ☐ Authentication fails consistently.
- ☐ The NameIdentifier claim cannot be the same as the claim URI.
- ☐ The NameIdentifier claim must be the same as the username stored in AD.

Licensed to Berend Post berend.post@it-university.nl

Question 61

A company is running a batch analysis every hour on their main transactional DB, running on an RDS MySQL instance, to populate their central Data Warehouse running on Redshift. During the execution of the batch, their transactional applications are very slow. When the batch completes they need to update the top management dashboard with the new data. The dashboard is produced by another system running on-premises that is currently started when a manually-sent email notifies that an update is required. The on-premises system cannot be modified because is managed by another team.

How would you optimize this scenario to solve performance issues and automate the process as much as possible?

- ☐ Replace RDS with Redshift for the batch analysis and SNS to notify the on-premises system to update the dashboard
- ☐ Replace RDS with Redshift for the oaten analysis and SQS to send a message to the on-premises system to update the dashboard
- ☐ Create an RDS Read Replica for the batch analysis and SNS to notify me on-premises system to update the dashboard
- ☐ Create an RDS Read Replica for the batch analysis and SQS to send a message to the on-premises system to update the dashboard.

Question 62

A user has suspended the scaling process on the Auto Scaling group. A scaling activity to increase the instance count was already in progress.

What effect will the suspension have on that activity?

- ☐ No effect. The scaling activity continues
- ☐ Pauses the instance launch and launches it only after Auto Scaling is resumed
- ☐ Terminates the instance
- ☐ Stops the instance temporary

Question 63

In which step of using AWS Direct Connect should the user determine the required port speed?

- ☐ Complete the Cross Connect
- ☐ Verify Your Virtual Interface
- ☐ Download Router Configuration
- ☐ Submit AWS Direct Connect Connection Request

Licensed to Berend Post berend.post@it-university.nl

Question 64

To serve Web traffic for a popular product your chief financial officer and IT director have purchased 10 m1.large heavy utilization Reserved Instances (RIs), evenly spread across two availability zones; Route 53 is used to deliver the traffic to an Elastic Load Balancer (ELB). After several months, the product grows even more popular and you need additional capacity. As a result, your company purchases two C3.2xlarge medium utilization Ris. You register the two c3.2xlarge instances with your ELB and quickly find that the m1.large instances are at 100% of capacity and the c3.2xlarge instances have significant capacity that's unused.

Which option is the most cost effective and uses EC2 capacity most effectively?

- ☐ Configure Autoscaling group and Launch Configuration with ELB to add up to 10 more on-demand m1.large instances when triggered by Cloudwatch. Shut off c3.2xlarge instances.
- ☐ Configure ELB with two c3.2xlarge instances and use on-demand Autoscaling group for up to two additional c3.2xlarge instances. Shut off m1.large instances.
- ☐ Route traffic to EC2 m1.large and c3.2xlarge instances directly using Route 53 latency based routing and health checks. Shut off ELB.
- ☐ Use a separate ELB for each instance type and distribute load to ELBs with Route 53 weighted round robin.

Question 65

A company is finalizing the architecture for its backup solution for applications running on AWS. All of the applications run on AWS and use at least two Availability Zones in each tier.

Company policy requires IT to durably store nightly backups f all its data in at least two locations: production and disaster recovery. The locations must be in different geographic regions. The company also needs the backup to be available to restore immediately at the production data center, and within 24 hours at the disaster recovery location. All backup processes must be fully automated.

What is the MOST cost-effective backup solution that will meet all requirements?

- ☐ Back up all the data to a large Amazon EBS volume attached to the backup media server in the production region. Run automated scripts to snapshot these volumes nightly, and copy these snapshots to the disaster recovery region.
- ☐ Back up all the data to Amazon S3 in the disaster recovery region. Use a lifecycle policy to move this data to Amazon Glacier in the production region immediately. Only the data is replicated; remove the data from the S3 bucket in the disaster recovery region.
- ☐ Back up all the data to Amazon Glacier in the production region. Set up cross-region replication of this data to Amazon Glacier in the disaster recovery region. Set up a lifecycle policy to delete any data older than 60 days.
- ☐ Back up all the data to Amazon S3 in the production region. Set up cross-region replication of this S3 bucket to another region and set up a lifecycle policy in the second region to immediately move this data to Amazon Glacier.

Licensed to Berend Post berend.post@it-university.nl

Question 66

A company runs a legacy system on a single m4.2xlarge Amazon EC2 instance with Amazon EBS2 storage. The EC2 instance runs both the web server and a self-managed Oracle database. A snapshot is made of the EBS volume every 12 hours, and an AMI was created from the fully configured EC2 instance.
A recent event that terminated the EC2 instance led to several hours of downtime. The application was successfully launched from the AMI, but the age of the EBS snapshot and the repair of the database resulted in the loss of 8 hours of data. The system was also down for 4 hours while the Systems Operators manually performed these processes.
What architectural changes will minimize downtime and reduce the chance of lost data?

☐ Create an Amazon CloudWatch alarm to automatically recover the instance. Create a script that will check and repair the database upon reboot. Subscribe the Operations team to the Amazon SNS message generated by the CloudWatch alarm.
☐ Run the application on m4.xlarge EC2 instances behind an Elastic Load Balancer/Application Load Balancer. Run the EC2 instances in an Auto Scaling group across multiple Availability Zones with a minimum instance count of two. Migrate the database to an Amazon RDS Oracle Multi-AZ DB instance.
☐ Run the application on m4.2xlarge EC2 instances behind an Elastic Load Balancer/Application Load Balancer. Run the EC2 instances in an Auto Scaling group access multiple Availability Zones with a minimum instance count of one. Migrate the database to an Amazon RDS Oracle Multi-AZ DB instance.
☐ Increase the web server instance count to two m4.xlarge instances and use Amazon Route 53 round-robin load balancing to spread the load. Enable Route 53 health checks on the web servers. Migrate the database to an Amazon RDS Oracle Multi-AZ DB instance.

Question 67

An organization is planning to use NoSQL DB for its scalable data needs. The organization wants to host an application securely in AWS VPC.

What action can be recommended to the organization?

☐ The organization should setup their own NoSQL cluster on the AWS instance and configure route tables and subnets.
☐ The organization should only use a DynamoDB because by default it is always a part of the default subnet provided by AWS.
☐ The organization should use a DynamoDB while creating a table within the public subnet.
☐ The organization should use a DynamoDB while creating a table within a private subnet.

Licensed to Berend Post berend.post@it-university.nl

Question 68

A user is creating a snapshot of an EBS volume. Which of the below statements is incorrect in relation to the creation of an EBS snapshot?

- ☐ Its incremental
- ☐ It is a point in time backup of the EBS volume
- ☐ It can be used to create an AMI
- ☐ It is stored in the same AZ as the volume

Question 69

In Amazon VPC, what is the default maximum number of BGP advertised routes allowed per route table?

- ☐ 15
- ☐ 100
- ☐ 5
- ☐ 10

Question 70

In the context of AWS CloudFormation, which of the following statements is correct?

- ☐ Actual resource names are a combination of the resource ID, stack, and logical resource name.
- ☐ Actual resource name is the stack resource name.
- ☐ Actual resource name is the logical resource name.
- ☐ Actual resource names are a combination of the stack and logical resource name.

Question 71

A user is trying to understand the detailed CloudWatch monitoring concept. Which of the below mentioned services does not provide detailed monitoring with CloudWatch?

- ☐ AWS RDS
- ☐ AWS ELB
- ☐ AWS Route53
- ☐ AWS EMR

Licensed to Berend Post berend.post@it-university.nl

Question 72

Your company runs a customer facing event registration site This site is built with a 3-tier architecture with web and application tier servers and a MySQL database The application requires 6 web tier servers and 6 application tier servers for normal operation, but can run on a minimum of 65% server capacity and a single MySQL database.

When deploying this application in a region with three availability zones (AZs) which architecture provides high availability?

☐ A web tier deployed across 2 AZs with 3 EC2 (Elastic Compute Cloud) instances in each AZ inside an Auto Scaling Group behind an ELB (elastic load balancer), and an application tier deployed across 2 AZs with 3 EC2 instances in each AZ inside an Auto Scaling Group behind an ELB and one RDS (Relational Database Service) instance deployed with read replicas in the other AZ.

☐ A web tier deployed across 3 AZs with 2 EC2 (Elastic Compute Cloud) instances in each AZ inside an Auto Scaling Group behind an ELB (elastic load balancer) and an application tier deployed across 3 AZs with 2 EC2 instances in each AZ inside an Auto Scaling Group behind an ELB and one RDS (Relational Database Service) Instance deployed with read replicas in the two other AZs.

☐ A web tier deployed across 2 AZs with 3 EC2 (Elastic Compute Cloud) instances in each AZ inside an Auto Scaling Group behind an ELB (elastic load balancer) and an application tier deployed across 2 AZs with 3 EC2 instances m each AZ inside an Auto Scaling Group behind an ELS and a Multi-AZ RDS (Relational Database Service) deployment.

☐ A web tier deployed across 3 AZs with 2 EC2 (Elastic Compute Cloud) instances in each AZ Inside an Auto Scaling Group behind an ELB (elastic load balancer). And an application tier deployed across 3 AZs with 2 EC2 instances in each AZ inside an Auto Scaling Group behind an ELB and a Multi-AZ RDS (Relational Database services) deployment.

Question 73

A company is migrating its on-premises build artifact server to an AWS solution. The current system consists of an Apache HTTP server that serves artifacts to clients on the local network, restricted by the perimeter firewall. The artifact consumers are largely build automation scripts that download artifacts via anonymous HTTP, which the company will be unable to modify within its migration timetable.

The company decides to move the solution to Amazon S3 static website hosting. The artifact consumers will be migrated to Amazon EC2 instances located within both public and private subnets in a virtual private cloud (VPC).

Which solution will permit the artifact consumers to download artifacts without modifying the existing automation scripts?

☐ Create a NAT gateway within a public subnet of the VPC. Add a default route pointing to the NAT gateway into the route table associated with the subnets containing consumers. Configure the bucket policy to allow the s3:ListBucket and s3:GetObject actions using the

Licensed to Berend Post berend.post@it-university.nl

condition IpAddress and the condition key aws:SourceIp matching the elastic IP address if the NAT gateway.

☐ Create a VPC endpoint and add it to the route table associated with subnets containing consumers. Configure the bucket policy to allow s3:ListBucket and s3:GetObject actions using the condition StringEquals and the condition key aws:sourceVpce matching the identification of the VPC endpoint.

☐ Create an IAM role and instance profile for Amazon EC2 and attach it to the instances that consume build artifacts. Configure the bucket policy to allow the s3:ListBucket and s3:GetObjects actions for the principal matching the IAM role created.

☐ Create a VPC endpoint and add it to the route table associated with subnets containing consumers. Configure the bucket policy to allow s3:ListBucket and s3:GetObject actions using the condition IpAddress and the condition key aws:SourceIp matching the VPC CIDR block.

Question 74

You need to develop and run some new applications on AWS and you know that Elastic Beanstalk and CloudFormation can both help as a deployment mechanism for a broad range of AWS resources.

Which of the following is TRUE statements when describing the differences between Elastic Beanstalk and CloudFormation?

☐ AWS Elastic Beanstalk introduces two concepts: The template, a JSON or YAML-format, text- based file

☐ Elastic Beanstalk supports AWS CloudFormation application environments as one of the AWS resource types.

☐ Elastic Beanstalk automates and simplifies the task of repeatedly and predictably creating groups of related resources that power your applications. CloudFormation does not.

☐ You can design and script custom resources in CloudFormation

Question 75

In the context of Amazon ElastiCache CLI, which of the following commands can you use to view all ElastiCache instance events for the past 24 hours?

☐ elasticache-events --duration 24

☐ elasticache-events --duration 1440

☐ elasticache-describe-events --duration 24

☐ elasticache describe-events --source-type cache-cluster --duration 1440

Licensed to Berend Post berend.post@it-university.nl

Question 76

In Amazon Elastic Compute Cloud, you can specify storage volumes in addition to the root device volume when you create an AMI or when launching a new instance using_____.

- ☐ block device mapping
- ☐ object mapping
- ☐ batch storage mapping
- ☐ datacenter mapping

Question 77

An organization is setting a website on the AWS VPC. The organization has blocked a few IPs to avoid a D-DOS attack.

How can the organization configure that a request from the above mentioned IPs does not access the application instances?

- ☐ Create an IAM policy for VPC which has a condition to disallow traffic from that IP address.
- ☐ Configure a security group at the subnet level which denies traffic from the selected IP.
- ☐ Configure the security group with the EC2 instance which denies access from that IP address.
- ☐ Configure an ACL at the subnet which denies the traffic from that IP address.

Question 78

You have custom Network File System (NFS) client settings for your Amazon Elastic File System (EFS). It takes up to three seconds for an Amazon Elastic Compute Cloud (EC2) instance to see a write operation performed on a file system from another Amazon EC2 instance.

Which of the following actions should you take to solve the custom NFS settings from causing delays in the write operation?

- ☐ Unmount and remount the file system with the noac option to disable attribute caching.
- ☐ Reduce the number of active users that have files open simultaneously on the instances.
- ☐ Verify that the IP address of the specified mount target is valid.
- ☐ Run the write operation from a different user ID on the same Amazon EC2 instance.

Licensed to Berend Post berend.post@it-university.nl

Question 79

A user is planning to host a Highly Available system on the AWS VPC. Which of the below mentioned statements is helpful in this scenario?

- ☐ Create VPC subnets in two separate availability zones and launch instances in different subnets.
- ☐ Create VPC with only one public subnet and launch instances in different AZs using that subnet.
- ☐ Create two VPCs in two separate zones and setup failover with ELB such that if one VPC fails it will divert traffic to another VPC.
- ☐ Create VPC with only one private subnet and launch instances in different AZs using that subnet.

Question 80

What types of identities do Amazon Cognito identity pools support?

- ☐ They support both authenticated and unauthenticated identities.
- ☐ They support only unauthenticated identities.
- ☐ They support neither authenticated nor unauthenticated identities.
- ☐ They support only authenticated identities.

Question 81

For Amazon EC2 issues, while troubleshooting AWS CloudFormation, you need to view the cloud-init and cfn logs for more information. Identify a directory to which these logs are published.

- ☐ /var/opt/log/ec2
- ☐ /var/log/lastlog
- ☐ /var/log/
- ☐ /var/log/ec2

Question 82

Does an AWS Direct Connect location provide access to Amazon Web Services in the region it is associated with as well as access to other US regions?

- ☐ No, it provides access only to the region it is associated with.
- ☐ No, it provides access only to the US regions other than the region it is associated with.
- ☐ Yes, it provides access.
- ☐ Yes, it provides access but only when there's just one Availability Zone in the region.

Licensed to Berend Post berend.post@it-university.nl

Question 83

A company runs a Windows Server host in a public subnet that is configured to allow a team of administrators to connect over RDP to troubleshoot issues with hosts in a private subnet. The host must be available at all times outside of a scheduled maintenance window, and needs to receive the latest operating system updates within 3 days of release.

What should be done to manage the host with the LEAST amount of administrative effort?

- ☐ Run the host in a single-instance AWS Elastic Beanstalk environment. Configure the environment with a custom AMI to use a hardened machine image from AWS Marketplace. Apply system updates with AWS Systems Manager Patch Manager.
- ☐ Run the host on AWS WorkSpaces. Use Amazon WorkSpaces Application Manager (WAM) to harden the host. Configure Windows automatic updates to occur every 3 days.
- ☐ Run the host in an Auto Scaling group with a minimum and maximum instance count of 1. Use a hardened machine image from AWS Marketplace. Apply system updates with AWS Systems Manager Patch Manager.
- ☐ Run the host in AWS OpsWorks Stacks. Use a Chief recipe to harden the AMI during instance launch. Use an AWS Lambda scheduled event to run the Upgrade Operating System stack command to apply system updates.

Question 84

Your department creates regular analytics reports from your company's log files All log data is collected in Amazon S3 and processed by daily Amazon Elastic MapReduce (EMR) jobs that generate daily PDF reports and aggregated tables in CSV format for an Amazon Redshift data warehouse. Your CFO requests that you optimize the cost structure for this system.

Which of the following alternatives will lower costs without compromising average performance of the system or data integrity for the raw data?

- ☐ Use reduced redundancy storage (RRS) for all data In S3. Use a combination of Spot Instances and Reserved Instances for Amazon EMR jobs. Use Reserved Instances for Amazon Redshift.
- ☐ Use reduced redundancy storage (RRS) for PDF and .csv data in S3. Add Spot Instances to EMR jobs. Use Spot Instances for Amazon Redshift.
- ☐ Use reduced redundancy storage (RRS) for PDF and .csv data In Amazon S3. Add Spot Instances to Amazon EMR jobs. Use Reserved Instances for Amazon Redshift.
- ☐ Use reduced redundancy storage (RRS) for all data in Amazon S3. Add Spot Instances to Amazon EMR jobs. Use Reserved Instances for Amazon Redshift.

Licensed to Berend Post berend.post@it-university.nl

Question 85

A company needs to run a software package that has a license that must be run on the same physical host for the duration of its use. The software package is only going to be used for 90 days. The company requires patching and restarting of all instances every 30 days.

How can these requirements be met using AWS?

- ☐ Run a dedicated instance with auto-placement disabled.
- ☐ Run the instance on a dedicated host with Host Affinity set to Host.
- ☐ Run an On-Demand instance with a Reserved Instance to ensure consistent placement.
- ☐ Run the instance on a licensed host with termination set for 90 days.

Question 86

An International company has deployed a multi-tier web application that relies on DynamoDB in a single region. For regulatory reasons they need disaster recovery capability in a separate region with a Recovery Time Objective of 2 hours and a Recovery Point Objective of 24 hours. They should synchronize their data on a regular basis and be able to provision me web application rapidly using CloudFormation.
The objective is to minimize changes to the existing web application, control the throughput of DynamoDB used for the synchronization of data and synchronize only the modified elements.

Which design would you choose to meet these requirements?

- ☐ Use AWS data Pipeline to schedule a DynamoDB cross region copy once a day, create a “Lastupdated” attribute in your DynamoDB table that would represent the timestamp of the last update and use it as a filter.
- ☐ Use EMR and write a custom script to retrieve data from DynamoDB in the current region using a SCAN operation and push it to DynamoDB in the second region.
- ☐ Use AWS data Pipeline to schedule an export of the DynamoDB table to S3 in the current region once a day then schedule another task immediately after it that will import data from S3 to DynamoDB in the other region.
- ☐ Send also each Ante into an SQS queue in me second region; use an auto-scaling group behind the SQS queue to replay the write in the second region.

Licensed to Berend Post berend.post@it-university.nl

Question 87

Someone is creating a VPC for their application hosting. He has created two private subnets in the same availability zone and created one subnet in a separate availability zone. He wants to make a High Availability system with an internal Elastic Load Balancer.

Which choice is true regarding internal ELBs in this scenario? (Choose 2 answers)

- ☐ Internal ELBs should only be launched within private subnets.
- ☐ Amazon ELB service does not allow subnet selection; instead it will automatically select all the available subnets of the VPC.
- ☐ Internal ELBs can support only one subnet in each availability zone.
- ☐ An internal ELB can support all the subnets irrespective of their zones.

Question 88

Identify a true statement about the statement ID (Sid) in IAM.

- ☐ You cannot expose the Sid in the IAM API.
- ☐ You cannot use a Sid value as a sub-ID for a policy document's ID for services provided by SQS and SNS.
- ☐ You can expose the Sid in the IAM API.
- ☐ You cannot assign a Sid value to each statement in a statement array.

Question 89

A hybrid network architecture must be used during a company's multi-year data center migration from multiple private data centers to AWS. The current data centers are linked together with private fiber. Due to unique legacy applications, NAT cannot be used. During the migration period, many applications will need access to other applications in both the data centers and AWS.

Which option offers a hybrid network architecture that is secure and highly available, that allows for high bandwidth and a multi-region deployment post-migration?

- ☐ Use AWS Direct Connect to each data center from different ISPs, and configure routing to failover to the other data center's Direct Connect if one fails. Ensure that no VPC CIDR blocks overlap one another or the on-premises network.
- ☐ Use multiple hardware VPN connections to AWS from the on-premises data center. Route different subnet traffic through different VPN connections. Ensure that no VPC CIDR blocks overlap one another or the on-premises network.
- ☐ Use a software VPN with clustering both in AWS and the on-premises data center, and route traffic through the cluster. Ensure that no VPC CIDR blocks overlap one another or the on-premises network.
- ☐ Use AWS Direct Connect and a VPN as backup, and configure both to use the same virtual private gateway and BGP. Ensure that no VPC CIDR blocks overlap one another or the on-premises network.

Licensed to Berend Post berend.post@it-university.nl

Question 90

Complete this statement: "When you load your table directly from an Amazon_____ table, you have the option to control the amount of provisioned throughput you consume."

- ☐ RDS
- ☐ DataPipeline
- ☐ DynamoDB
- ☐ S3

Question 91

Can a Direct Connect link be connected directly to the Internet?

- ☐ Yes, this can be done if you pay for it.
- ☐ Yes, this can be done only for certain regions.
- ☐ Yes
- ☐ No

Question 92

In Amazon Redshift, how many slices does a dw2.8xlarge node have?

- ☐ 16
- ☐ 8
- ☐ 32
- ☐ 2

Question 93

A sys admin is maintaining an application on AWS. The application is installed on EC2 and user has configured ELB and Auto Scaling. Considering future load increase, the user is planning to launch new servers proactively so that they get registered with ELB.

How can the user add these instances with Auto Scaling?

- ☐ Decrease the minimum limit of the Auto Scaling group
- ☐ Increase the maximum limit of the Auto Scaling group
- ☐ Launch an instance manually and register it with ELB on the fly
- ☐ Increase the desired capacity of the Auto Scaling group

Licensed to Berend Post berend.post@it-university.nl

Question 94

Within an IAM policy, can you add an IfExists condition at the end of a Null condition?

- ☐ Yes, you can add an IfExists condition at the end of a Null condition but not in all Regions.
- ☐ Yes, you can add an IfExists condition at the end of a Null condition depending on the condition.
- ☐ No, you cannot add an IfExists condition at the end of a Null condition.
- ☐ Yes, you can add an IfExists condition at the end of a Null condition.

Question 95

A user has launched two EBS backed EC2 instances in the US-East-1a region. The user wants to change the zone of one of the instances. How can the user change it?

- ☐ It is not possible to change the zone of an instance after it is launched
- ☐ From the AWS EC2 console, select the Actions - > Change zones and specify the new zone
- ☐ The zone can only be modified using the AWS CLI
- ☐ Stop one of the instances and change the availability zone

Question 96

IAM Secure and Scalable is an organization which provides scalable and secure SAAS to its clients. They are planning to host a web server and App server on AWS VPC as separate tiers. The organization wants to implement the scalability by configuring Auto Scaling and load balancer with their app servers (middle tier) too.

Which of the below mentioned options suits their requirements?

- ☐ Since ELB is internet facing, it is recommended to setup HAProxy as the Load balancer within the VPC.
- ☐ Create an Internet facing ELB with VPC and configure all the App servers with it.
- ☐ The user should make ELB with EC2-CLASSIC and enable SSH with it for security.
- ☐ Create an Internal Load balancer with VPC and register all the App servers with it.

Question 97

A user has launched an EBS optimized instance with EC2. Which of the below mentioned options is the correct statement?

- ☐ It provides additional dedicated capacity for EBS IO
- ☐ The attached EBS will have greater storage capacity
- ☐ The user will have a PIOPS based EBS volume
- ☐ It will be launched on dedicated hardware in VPC

Licensed to Berend Post berend.post@it-university.nl

Question 98

How does AWS Data Pipeline execute activities on on-premise resources or AWS resources that you manage?

- ☐ By supplying a Task Runner package that can be installed on your on-premise hosts
- ☐ None of these
- ☐ By supplying a Task Runner file that the resources can access for execution
- ☐ By supplying a Task Runner json script that can be installed on your on-premise hosts

Question 99

A company wants to ensure that the workloads for each of its business units have complete autonomy and a minimal blast radius in AWS. The Security team must be able to control access to the resources and services in the account to ensure that particular services are not used by the business units.

How can a Solutions Architect achieve the isolation requirements?

- ☐ Create individual accounts for each business unit and add the account to an OU in AWS Organizations. Modify the OU to ensure that the particular services are blocked. Federate each account with an IdP, and create separate roles for the business units and the Security team.
- ☐ Create individual accounts for each business unit. Federate each account with an IdP and create separate roles and policies for business units and the Security team.
- ☐ Create one shared account for the entire company. Create separate VPCs for each business unit. Create individual IAM policies and resource tags for each business unit. Federate each account with an IdP, and create separate roles for the business units and the Security team.
- ☐ Create one shared account for the entire company. Create individual IAM policies and resource tags for each business unit. Federate the account with an IdP, and create separate roles for the business units and the Security team.

Question 100

A user is planning to launch multiple EC2 instance same as current running instance.

Which of the below mentioned parameters is not copied by Amazon EC2 in the launch wizard when the user has selected the option "Launch more like this"?

- ☐ Termination protection
- ☐ Tenancy setting
- ☐ Storage
- ☐ Shutdown behavior

Licensed to Berend Post berend.post@it-university.nl

Question 101

A company is running a .NET three-tier web application on AWS. The team currently uses XL storage optimized instances to store serve the website's image and video files on local instance storage. The company has encountered issues with data loss from replication and instance failures. The Solutions Architect has been asked to redesign this application to improve its reliability while keeping costs low.

Which solution will meet these requirements?

- ☐ Set up a new Amazon EFS share, move all image and video files to this share, and then attach this new drive as a mount point to all existing servers. Create an Elastic Load Balancer with Auto Scaling general purpose instances. Enable Amazon CloudFront to the Elastic Load Balancer. Enable Cost Explorer and use AWS Trusted advisor checks to continue monitoring the environment for future savings.
- ☐ Implement Auto Scaling with general purpose instance types and an Elastic Load Balancer. Enable an Amazon CloudFront distribution to Amazon S3 and move images and video files to Amazon S3. Reserve general purpose instances to meet base performance requirements. Use Cost Explorer and AWS Trusted Advisor checks to continue monitoring the environment for future savings.
- ☐ Move the entire website to Amazon S3 using the S3 website hosting feature. Remove all the web servers and have Amazon S3 communicate directly with the application servers in Amazon VPC.
- ☐ Use AWS Elastic Beanstalk to deploy the .NET application. Move all images and video files to Amazon EFS. Create an Amazon CloudFront distribution that points to the EFS share. Reserve the m4.4xl instances needed to meet base performance requirements.

Question 102

A user has created a VPC with public and private subnets using the VPC Wizard. The VPC has CIDR 20.0.0.0/16. The private subnet uses CIDR 20.0.0.0/24.

Which of the below mentioned entries are required in the main route table to allow the instances in VPC to communicate with each other?

- ☐ Destination : 20.0.0.0/0 and Target : ALL
- ☐ Destination : 20.0.0.0/16 and Target : Local
- ☐ Destination : 20.0.0.0/24 and Target : Local
- ☐ Destination : 20.0.0.0/16 and Target : ALL

Licensed to Berend Post berend.post@it-university.nl

Question 103

A user is using CloudFormation to launch an EC2 instance and then configure an application after the instance is launched. The user wants the stack creation of ELB and AutoScaling to wait until the EC2 instance is launched and configured properly.

How can the user configure this?

- ☐ The user can use the DependentCondition resource to hold the creation of the other dependent resources.
- ☐ It is not possible that the stack creation will wait until one service is created and launched.
- ☐ The user can use the HoldCondition resource to wait for the creation of the other dependent resources.
- ☐ The user can use the WaitCondition resource to hold the creation of the other dependent resources.

Licensed to Berend Post berend.post@it-university.nl

Question 104

A Solutions Architect must design a highly available, stateless, REST service. The service will require multiple persistent storage layers for service object meta information and the delivery of content. Each request needs to be authenticated and securely processed. There is a requirement to keep costs as low as possible?

How can these requirements be met?

- ☐ Use AWS Fargate to host a container that runs a self-contained REST service. Set up an Amazon ECS service that is fronted by an Application Load Balancer (ALB). Use a custom authenticator to control access to the API. Store request meta information in Amazon DynamoDB with Auto Scaling and static content in a secured S3 bucket. Make secure signed requests for Amazon S3 objects and proxy the data through the REST service interface.
- ☐ Use AWS Fargate to host a container that runs a self-contained REST service. Set up an ECS service that is fronted by a cross-zone ALB. Use an Amazon Cognito user pool to control access to the API. Store request meta information in DynamoDB with Auto Scaling and static content in a secured S3 bucket. Generate presigned URLs when returning references to content stored in Amazon S3.
- ☐ Set up Amazon API Gateway and create the required API resources and methods. Use an Amazon Cognito user pool to control access to the API. Configure the methods to use AWS Lambda proxy integrations, and process each resource with a unique AWS Lambda function. Store request meta information in DynamoDB with Auto Scaling and static content in a secured S3 bucket. Generate presigned URLs when returning references to content stored in Amazon S3.
- ☐ Set up Amazon API Gateway and create the required API resources and methods. Use an Amazon API Gateway custom authorizer to control access to the API. Configure the methods to use AWS Lambda custom integrations, and process each resource with a unique Lambda function. Store request meta information in an Amazon ElastiCache Multi-AZ cluster and static content in a secured S3 bucket. Generate presigned URLs when returning references to content stored in Amazon S3.

Question 105

A user has created a MySQL RDS instance with PIOPS. Which of the below mentioned statements will help user understand the advantage of PIOPS?

- ☐ The user can achieve additional dedicated capacity for the EBS I/O with an enhanced RDS option
- ☐ It uses a standard EBS volume with optimized configuration the stacks
- ☐ It uses optimized EBS volumes and optimized configuration stacks
- ☐ It provides a dedicated network bandwidth between EBS and RDS

Licensed to Berend Post berend.post@it-university.nl

Question 106

A large company is migrating its entire IT portfolio to AWS. Each business unit in the company has a standalone AWS account that supports both development and test environments. New accounts to support production workloads will be needed soon.

The Finance department requires a centralized method for payment but must maintain visibility into each group's spending to allocate costs.

The Security team requires a centralized mechanism to control IAM usage in all the company's accounts.

What combination of the following options meet the company's needs with LEAST effort? (Choose two.)

- ☐ Use a collection of parameterized AWS CloudFormation templates defining common IAM permissions that are launched into each account. Require all new and existing accounts to launch the appropriate stacks to enforce the least privilege model.
- ☐ Use AWS Organizations to create a new organization from a chosen payer account and define an organizational unit hierarchy. Invite the existing accounts to join the organization and create new accounts using Organizations.
- ☐ Require each business unit to use its own AWS accounts. Tag each AWS account appropriately and enable Cost Explorer to administer chargebacks.
- ☐ Enable all features of AWS Organizations and establish appropriate service control policies that filter IAM permissions for sub-accounts.
- ☐ Consolidate all of the company's AWS accounts into a single AWS account. Use tags for billing purposes and IAM's Access Advice feature to enforce the least privilege model.

Question 107

What is the average queue length recommended by AWS to achieve a lower latency for the 200 PIOPS EBS volume?

- ☐ 5
- ☐ 1
- ☐ 2
- ☐ 4

Licensed to Berend Post berend.post@it-university.nl

Question 108

A Solutions Architect has created an AWS CloudFormation template for a three-tier application that contains an Auto Scaling group of Amazon EC2 instances running a custom AMI.

The Solutions Architect wants to ensure that future updates to the custom AMI can be deployed to a running stack by first updating the template to refer to the new AMI, and then invoking UpdateStack to replace the EC2 instances with instances launched from the new AMI.

How can updates to the AMI be deployed to meet these requirements?

- ☐ Create a change set for a new version of the template, view the changes to the running EC2 instances to ensure that the AMI is correctly updated, and then execute the change set.
- ☐ Edit the AWS::AutoScaling: :LaunchConfiguration resource in the template, changing its DeletionPolicy to Replace.
- ☐ Edit the AWS::AutoScaling: :LaunchConfiguration resource in the template, inserting an UpdatePolicy attribute.
- ☐ Create a new stack from the updated template. Once it is successfully deployed, modify the DNS records to point to the new stack and delete the old stack.

Question 109

A Solutions Architect must establish a patching plan for a large mixed fleet of Windows and Linux servers. The patching plan must be implemented securely, be audit ready, and comply with the company's business requirements.

Which option will meet these requirements with MINIMAL effort?

- ☐ Install and use an OS-native patching service to manage the update frequency and release approval for all instances. Use AWS Config to verify the OS state on each instance and report on any patch compliance issues.
- ☐ Use AWS Systems Manager on all instances to manage patching. Test patches outside of production and then deploy during a maintenance window with the appropriate approval.
- ☐ Use AWS OpsWorks for Chef Automate to run a set of scripts that will iterate through all instances of a given type. Issue the appropriate OS command to get and install updates on each instance, including any required restarts during the maintenance window.
- ☐ Migrate all applications to AWS OpsWorks and use OpsWorks automatic patching support to keep the OS up-to-date following the initial installation. Use AWS Config to provide audit and compliance reporting.

Licensed to Berend Post berend.post@it-university.nl

Question 110

When you put objects in Amazon S3, what is the indication that an object was successfully stored?

- ☐ A HTTP 200 result code and MD5 checksum, taken together, indicate that the operation was successful.
- ☐ Amazon S3 is engineered for 99.999999999% durability. Therefore there is no need to confirm that data was inserted.
- ☐ A success code is inserted into the S3 object metadata.
- ☐ Each S3 account has a special bucket named _s3_logs. Success codes are written to this bucket with a timestamp and checksum.

Question 111

You are playing around with setting up stacks using JSON templates in CloudFormation to try and understand them a little better. You have set up about 5 or 6 but now start to wonder if you are being charged for these stacks.

What is AWS's billing policy regarding stack resources?

- ☐ You are not charged for the stack resources if they are not taking any traffic.
- ☐ You are charged for the stack resources for the time they were operating (but not if you deleted the stack within 30 minutes)
- ☐ You are charged for the stack resources for the time they were operating (but not if you deleted the stack within 60 minutes)
- ☐ You are charged for the stack resources for the time they were operating (even if you deleted the stack right away)

Question 112

A Provisioned IOPS volume must be at least _____ GB in size:

- ☐ 20
- ☐ 10
- ☐ 50
- ☐ 1

Question 113

A user is configuring MySQL RDS with PIOPS. What should be the minimum PIOPS that the user should provision?

- ☐ 1000
- ☐ 200
- ☐ 2000
- ☐ 500

Licensed to Berend Post berend.post@it-university.nl

Question 114

What is the maximum length for a certificate ID in AWS IAM?

- ☐ 1024 characters
- ☐ 512 characters
- ☐ 64 characters
- ☐ 128 characters

Question 115

A company has developed a new billing application that will be released in two weeks. Developers are testing the application running on 10 EC2 instances managed by an Auto Scaling group in subnet 172.31.0.0/24 within VPC A with CIDR block 172.31.0.0/16. The Developers noticed connection timeout errors in the application logs while connecting to an Oracle database running on an Amazon EC2 instance in the same region within VPC B with CIDR block 172.50.0.0/16. The IP of the database instance is hard-coded in the application instances.

Which recommendations should a Solutions Architect present to the Developers to solve the problem in a secure way with minimal maintenance and overhead?

- ☐ Disable the SrcDestCheck attribute for all instances running the application and Oracle Database. Change the default route of VPC A to point ENI of the Oracle Database that has an IP address assigned within the range of 172.50.0.0/26
- ☐ Create and attach internet gateways for both VPCs. Configure default routes to the Internet gateways for both VPCs. Assign an Elastic IP for each Amazon EC2 instance in VPC A
- ☐ Create a VPC peering connection between the two VPCs and add a route to the routing table of VPC A that points to the IP address range of 172.50.0.0/16
- ☐ Create an additional Amazon EC2 instance for each VPC as a customer gateway; create one virtual private gateway (VGW) for each VPC, configure an end-to-end VPC, and advertise the routes for 172.50.0.0/16

Question 116

A user is trying to create a PIOPS EBS volume with 4000 IOPS and 100 GB size. AWS does not allow the user to create this volume.

What is the possible root cause for this?

- ☐ PIOPS is supported for EBS higher than 500 GB size
- ☐ The maximum IOPS supported by EBS is 3000
- ☐ The ratio between IOPS and the EBS volume is higher than 30
- ☐ The ratio between IOPS and the EBS volume is lower than 50

Licensed to Berend Post berend.post@it-university.nl

Question 117

A group of research institutions and hospitals are in a partnership to study 2 PBs of genomic data. The institute that owns the data stores it in an Amazon S3 bucket and updates it regularly. The institute would like to give all of the organizations in the partnership read access to the data. All members of the partnership are extremely cost-conscious, and the institute that owns the account with the S3 bucket is concerned about covering the costs for requests and data transfers from Amazon S3.

Which solution allows for secure datasharing without causing the institute that owns the bucket to assume all the costs for S3 requests and data transfers?

- ☐ Ensure that all organizations in the partnership have AWS accounts. In the account with the S3 bucket, create a cross-account role for each account in the partnership that allows read access to the data. Have the organizations assume and use that read role when accessing the data.
- ☐ Ensure that all organizations in the partnership have AWS accounts. Create a bucket policy on the bucket that owns the data. The policy should allow the accounts in the partnership read access to the bucket. Enable Requester Pays on the bucket. Have the organizations use their AWS credentials when accessing the data.
- ☐ Ensure that all organizations in the partnership have AWS accounts. Configure buckets in each of the accounts with a bucket policy that allows the institute that owns the data the ability to write to the bucket. Periodically sync the data from the institute's account to the other organizations. Have the organizations use their AWS credentials when accessing the data using their accounts.
- ☐ Ensure that all organizations in the partnership have AWS accounts. In the account with the S3 bucket, create a cross-account role for each account in the partnership that allows read access to the data. Enable Requester Pays on the bucket. Have the organizations assume and use that read role when accessing the data.

Question 118

You create an Amazon Elastic File System (EFS) file system and mount targets for the file system in your Virtual Private Cloud (VPC). Identify the initial permissions you can grant to the group root of your file system.

- ☐ write-execute-modify
- ☐ read-execute
- ☐ read-write-modify
- ☐ read-write

Licensed to Berend Post berend.post@it-university.nl

Question 119

An organization is planning to host a web application in the AWS VPC. The organization does not want to host a database in the public cloud due to statutory requirements.

How can the organization setup in this scenario?

☐ The organization should plan the app server on the public subnet and database in the organization's data center and connect them with the VPN gateway.
☐ The organization should plan the app server on the public subnet and use RDS with the private subnet for a secure data operation.
☐ The organization should use the public subnet for the app server and use RDS with a storage gateway to access as well as sync the data securely from the local data center.
☐ The organization should plan the app server on the public subnet and database in a private subnet so it will not be in the public cloud.

Question 120

A company runs a video processing platform. Files are uploaded by users who connect to a web server, which stores them on an Amazon EFS share. This web server is running on a single Amazon EC2 instance. A different group of instances, running in an Auto Scaling group, scans the EFS share directory structure for new files to process and generates new videos (thumbnails, different resolution, compression, etc.) according to the instructions file, which is uploaded along with the video files. A different application running on a group of instances managed by an Auto Scaling group processes the video files and then deletes them from the EFS share. The results are stored in an S3 bucket. Links to the processed video files are emailed to the customer.

The company has recently discovered that as they add more instances to the Auto Scaling Group, many files are processed twice, so image processing speed is not improved. The maximum size of these video files is 2GB.

What should the Solutions Architect do to improve reliability and reduce the redundant processing of video files?

☐ Modify the web application to upload the video files directly to Amazon S3. Use Amazon CloudWatch Events to trigger an AWS Lambda function every time a file is uploaded, and have this Lambda function put a message into an Amazon SQS queue. Modify the video processing application to read from SQS queue for new files and use the queue depth metric to scale instances in the video processing Auto Scaling group.
☐ Set up a cron job on the web server instance to synchronize the contents of the EFS share into Amazon S3. Trigger an AWS Lambda function every time a file is uploaded to process the video file and store the results in Amazon S3. Using Amazon CloudWatch Events trigger an Amazon SES job to send an email to the customer containing the link to the processed file.
☐ Rewrite the web application to run directly from Amazon S3 and use Amazon API Gateway to upload the video files to an S3 bucket. Use an S3 trigger to run an AWS Lambda function each time a file is uploaded to process and store new video files in a different

Licensed to Berend Post berend.post@it-university.nl

bucket. Using CloudWatch Events, trigger an SES job to send an email to the customer containing the link to the processed file.

☐ Rewrite the application to run from Amazon S3 and upload the video files to an S3 bucket. Each time a new file is uploaded, trigger an AWS Lambda function to put a message in an SQS queue containing the link and the instructions. Modify the video processing application to read from the SQS queue and the S3 bucket. Use the queue depth metric to adjust the size of the Auto Scaling group for video processing instances.

Question 121

You have set up Auto Scaling to automatically scale in. Consequently, you must decide which instances Auto Scaling should end first.

What should you use to configure this?

☐ An Elastic Load Balancer
☐ A termination policy
☐ An IAM role
☐ Another scaling group

Question 122

Once the user has set ElastiCache for an application and it is up and running, which services, does Amazon not provide for the user:

☐ The ability for client programs to automatically identify all of the nodes in a cache cluster, and to initiate and maintain connections to all of these nodes

☐ Automating common administrative tasks such as failure detection and recovery, and software patching.

☐ Providing default Time to Live (TTL) in the AWS Elasticache Redis Implementation for different type of data.

☐ Providing detailed monitoring metrics associated with your Cache Nodes, enabling you to diagnose and react to issues very quickly

Licensed to Berend Post berend.post@it-university.nl

Question 123

An internal security audit of AWS resources within a company found that a number of Amazon EC2 instances running Microsoft Windows workloads were missing several important operating system-level patches. A Solutions Architect has been asked to fix existing patch deficiencies, and to develop a workflow to ensure that future patching requirements are identified and taken care of quickly. The Solutions Architect has decided to use AWS Systems Manager. It is important that EC2 instance reboots do not occur at the same time on all Windows workloads to meet organizational uptime requirements.

Which workflow will meet these requirements in an automated manner?

☐ Add a Patch Group tag with a value of Windows Servers to all existing EC2 instances. Ensure that all Windows EC2 instances are assigned this tag. Associate the AWS-DefaultPatchBaseline to the Windows servers patch group. Define an AWS Systems Manager maintenance window, conduct patching within it, and associate it with the Windows Servers patch group. Register instances with the maintenance window using associated subnet IDs. Assign the AWS-RunPatchBaseline document as a task within each maintenance window.

☐ Add a Patch Group tag a value of Windows Servers to all existing EC2 instances. Ensure that all Windows EC2 instances are assigned this tag. Associate the AWS-WindowsPatchBaseline document as a task associated with the Windows Servers patch group. Create an Amazon CloudWatch Events rule configured to use a cron expression to schedule the execution of patching using the AWS Systems Manager run command. Create an AWS Systems Manager State Manager document to define commands to be executed during patch execution.

☐ Add a Patch Group tag with a value of either Windows Servers1 or Windows Server2 to all existing EC2 instances. Ensure that all Windows EC2 instances are assigned this tag. Associate the AWS-DefaultPatchBaseline with both Windows Servers patch groups. Define two non-overlapping AWS Systems Manager maintenance windows, conduct patching within them, and associate each with a different patch group. Register targets with specific maintenance windows using the Patch Group tags. Assign the AWS-RunPatchBaseline document as a task within each maintenance window.

☐ Add a Patch Group tag with a value of either Windows servers1 or Windows Server2 to all existing EC2 instances. Ensure that all Windows EC2 instances are assigned this tag. Associate the AWS-WindowsPatchBaseline with both Windows Servers patch groups. Define two non-overlapping AWS Systems Manager maintenance windows, conduct patching within them, and associate each with a different patch group. Assign the AWS-RunWindowsPatchBaseline document as a task within each maintenance window. Create an AWS Systems Manager State Manager document to define commands to be executed during patch execution.

Licensed to Berend Post berend.post@it-university.nl

Question 124

A user has created a VPC with CIDR 20.0.0.0/16. The user has created one subnet with CIDR 20.0.0.0/16 in this VPC. The user is trying to create another subnet with the same VPC for CIDR 20.0.0.1/24.

What will happen in this scenario?

- ☐ The VPC will modify the first subnet CIDR automatically to allow the second subnet IP range
- ☐ The second subnet will be created
- ☐ It will throw a CIDR overlaps error
- ☐ It is not possible to create a subnet with the same CIDR as VPC

Question 125

An organization has setup RDS with VPC. The organization wants RDS to be accessible from the internet. Which of the below mentioned configurations is not required in this scenario?

- ☐ The organization must enable the parameter in the console which makes the RDS instance publicly accessible.
- ☐ The organization must allow access from the internet in the RDS VPC security group,
- ☐ The organization must setup RDS with the subnet group which has an external IP.
- ☐ The organization must enable the VPC attributes DNS hostnames and DNS resolution.

Question 126

A user is creating a Provisioned IOPS volume. What is the maximum ratio the user should configure between Provisioned IOPS and the volume size?

- ☐ 30 to 1
- ☐ 50 to 1
- ☐ 10 to 1
- ☐ 20 to 1

Question 127

You are setting up some EBS volumes for a customer who has requested a setup which includes a RAID (redundant array of inexpensive disks). AWS has some recommendations for RAID setups.

Which RAID setup is not recommended for Amazon EBS?

- ☐ RAID 1 only
- ☐ RAID 5 only
- ☐ RAID 5 and RAID 6
- ☐ RAID 0 only

Licensed to Berend Post berend.post@it-university.nl

Question 128

You are running a successful multitier web application on AWS and your marketing department has asked you to add a reporting tier to the application. The reporting tier will aggregate and publish status reports every 30 minutes from user-generated information that is being stored in your web application s database. You are currently running a Multi-AZ RDS MySQL instance for the database tier. You also have implemented Elasticache as a database caching layer between the application tier and database tier.

Please select the answer that will allow you to successfully implement the reporting tier with as little impact as possible to your database.

- ☐ Continually send transaction logs from your master database to an S3 bucket and generate the reports off the S3 bucket using S3 byte range requests.
- ☐ Generate the reports by querying the synchronously replicated standby RDS MySQL instance maintained through Multi-AZ.
- ☐ Launch a RDS Read Replica connected to your Multi AZ master database and generate reports by querying the Read Replica.
- ☐ Generate the reports by querying the ElastiCache database caching tier.

Question 129

An organization has two Amazon EC2 instances:
The first is running an ordering application and an inventory application.
The second is running a queuing system.

During certain times of the year, several thousand orders are placed per second. Some orders were lost when the queuing system was down. Also, the organization's inventory application has the incorrect quantity of products because some orders were processed twice.

What should be done to ensure that the applications can handle the increasing number of orders?

- ☐ Put the ordering and inventory applications into their own AWS Lambda functions. Have the ordering application write the messages into an Amazon SQS FIFO queue.
- ☐ Put the ordering and inventory applications into their own Amazon ECS containers and create an Auto Scaling group for each application. Then, deploy the message queuing server in multiple Availability Zones.
- ☐ Put the ordering and inventory applications into their own Amazon EC2 instances, and create an Auto Scaling group for each application. Use Amazon SQS standard queues for the incoming orders, and implement idempotency in the inventory application.
- ☐ Put the ordering and inventory applications into their own Amazon EC2 instances. Write the incoming orders to an Amazon Kinesis data stream Configure AWS Lambda to poll the stream and update the inventory application.

Licensed to Berend Post berend.post@it-university.nl

Question 130

A company is moving a business-critical application onto AWS. It is a traditional three-tier web application using an Oracle database. Data must be encrypted in transit and at rest. The database hosts 12 TB of data. Network connectivity to the source Oracle database over the internal is allowed, and the company wants to reduce the operational costs by using AWS Managed Services where possible. All primary keys only; however, it contains many Binary Large Object (BLOB) fields. It was not possible to use the database's native replication tools because of licensing restrictions.

Which database migration solution will result in the LEAST amount of impact to the application's availability?

☐ Provision an Amazon RDS for Oracle instance. Host the RDS database within a virtual private cloud (VPC) subnet with internet access, and set up the RDS database as an encrypted Read Replica of the source database. Use SSL to encrypt the connection between the two databases. Monitor the replication performance by watching the RDS ReplicaLag metric. During the application maintenance window, shut down the on-premises database and switch over the application connection to the RDS instance when there is no more replication lag. Promote the Read Replica into a standalone database instance.

☐ Provision an Amazon EC2 instance and install the same Oracle database software. Create a backup of the source database using the supported tools. During the application maintenance window, restore the backup into the Oracle database running in the EC2 instance. Set up an Amazon RDS for Oracle instance, and create an import job between the database hosted in AWS. Shut down the source database and switch over the database connections to the RDS instance when the job is complete.

☐ Use AWS DMS to load and replicate the dataset between the on-premises Oracle database and the replication instance hosted on AWS. Provision an Amazon RDS for Oracle instance with Transparent Data Encryption (TDE) enabled and configure it as target for the replication instance. Create a customer-managed AWS KMS master key to set it as the encryption key for the replication instance. Use AWS DMS tasks to load the data into the target RDS instance. During the application maintenance window and after the load tasks reach the ongoing replication phase, switch the database connections to the new database.

☐ Create a compressed full database backup on the on-premises Oracle database during an application maintenance window. While the backup is being performed, provision a 10 Gbps AWS Direct Connect connection to increase the transfer speed of the database backup files to Amazon S3, and shorten the maintenance window period. Use SSL/TLS to copy the files over the Direct Connect connection. When the backup files are successfully copied, start the maintenance window, and rise any of the Amazon RDS supported tools to import the data into a newly provisioned Amazon RDS for Oracle instance with encryption enabled. Wait until the data is fully loaded and switch over the database connections to the new database. Delete the Direct Connect connection to cut unnecessary charges.

Licensed to Berend Post berend.post@it-university.nl

Question 131

A company currently uses Amazon EBS and Amazon RDS for storage purposes. The company intends to use a pilot light approach for disaster recovery in a different AWS Region. The company has an RTO of 6 hours and an RPO of 24 hours.

Which solution would achieve the requirements with MINIMAL cost?

- ☐ Use AWS Lambda to create daily EBS and RDS snapshots, and copy them to the disaster recovery region. Use Amazon Route 53 with active-passive failover configuration. Use Amazon EC2 in an Auto Scaling group with the capacity set to 0 in the disaster recovery region.
- ☐ Use AWS Lambda to create daily EBS and RDS snapshots, and copy them to the disaster recovery region. Use Amazon Route 53 with active-active failover configuration. Use Amazon EC2 in an Auto Scaling group configured in the same way as in the primary region.
- ☐ Use Amazon ECS to handle long-running tasks to create daily EBS and RDS snapshots, and copy to the disaster recovery region. Use Amazon Route 53 with active-passive failover configuration. Use Amazon EC2 in an Auto Scaling group with the capacity set to 0 in the disaster recovery region.
- ☐ Use EBS and RDS cross-region snapshot copy capability to create snapshots in the disaster recovery region. Use Amazon Route 53 with active-active failover configuration. Use Amazon EC2 in an Auto Scaling group with the capacity set to 0 in the disaster recovery region.

Question 132

A company uses Amazon S3 to store documents that may only be accessible to an Amazon EC2 instance in a certain virtual private cloud (VPC). The company fears that a malicious insider with access to this instance could also set up an EC2 instance in another VPC to access these documents.

Which of the following solutions will provide the required protection?

- ☐ Use an S3 VPC endpoint and an S3 bucket policy to limit access to this VPC endpoint.
- ☐ Use EC2 instance profiles and an S3 bucket policy to limit access to the role attached to the instance profile.
- ☐ Use S3 client-side encryption and store the key in the instance metadata.
- ☐ Use S3 server-side encryption and protect the key with an encryption context.

Question 133

Regarding Identity and Access Management (IAM), Which type of special account belonging to your application allows your code to access Google services programmatically?

- ☐ Service account
- ☐ Simple Key
- ☐ OAuth
- ☐ Code account

Licensed to Berend Post berend.post@it-university.nl

Question 134

Which of the following is true of Amazon EBS encryption keys?

- ☐ Amazon EBS encryption uses the Customer Master Key (CMK) to create an AWS Key Management Service (AWS KMS) master key.
- ☐ Amazon EBS encryption uses the EBS Magnetic key to create an AWS Key Management Service (AWS KMS) master key.
- ☐ Amazon EBS encryption uses the EBS Magnetic key to create a Customer Master Key (CMK).
- ☐ Amazon EBS encryption uses the AWS Key Management Service (AWS KMS) master key to create a Customer Master Key (CMK).

Question 135

A user has created an AWS AMI. The user wants the AMI to be available only to his friend and not anyone else. How can the user manage this?

- ☐ Share the AMI with the community and setup the approval workflow before anyone launches it.
- ☐ It is not possible to share the AMI with the selected user.
- ☐ Share the AMI with a friend's AWS account ID.
- ☐ Share the AMI with a friend's AWS login ID.

Question 136

An online retailer needs to regularly process large product catalogs, which are handled in batches. These are sent out to be processed by people using the Amazon Mechanical Turk service, but the retailer has asked its Solutions Architect to design a workflow orchestration system that allows it to handle multiple concurrent Mechanical Turk operations, deal with the result assessment process, and reprocess failures.

Which of the following options gives the retailer the ability to interrogate the state of every workflow with the LEAST amount of implementation effort?

- ☐ Trigger Amazon CloudWatch alarms based upon message visibility in multiple Amazon SQS queues (one queue per workflow stage) and send messages via Amazon SNS to trigger AWS Lambda functions to process the next step. Use Amazon ES and Kibana to visualize Lambda processing logs to see the workflow states.
- ☐ Hold workflow information in an Amazon RDS instance with AWS Lambda functions polling RDS for status changes. Worker Lambda functions then process the next workflow steps. Amazon QuickSight will visualize workflow states directly out of Amazon RDS.
- ☐ Build the workflow in AWS Step Functions, using it to orchestrate multiple concurrent workflows. The status of each workflow can be visualized in the AWS Management Console, and historical data can be written to Amazon S3 and visualized using Amazon QuickSight.
- ☐ Use Amazon SWF to create a workflow that handles a single batch of catalog records with multiple worker tasks to extract the data, transform it, and send it through Mechanical

Licensed to Berend Post berend.post@it-university.nl

Turk. Use Amazon ES and Kibana to visualize AWS Lambda processing logs to see the workflow states.

Question 137

In the context of AWS Cloud Hardware Security Module(HSM), does your application need to reside in the same VPC as the CloudHSM instance?

- ☐ No, but the server or instance on which your application and the HSM client is running must have network (IP) reachability to the HSM.
- ☐ Yes, always
- ☐ No, but they must reside in the same Availability Zone.
- ☐ No, but it should reside in same Availability Zone as the DB instance.

Question 138

With respect to AWS Lambda permissions model, at the time you create a Lambda function, you specify an IAM role that AWS Lambda can assume to execute your Lambda function on your behalf. This role is also referred to as the_____role.

- ☐ configuration
- ☐ execution
- ☐ delegation
- ☐ dependency

Question 139

In AWS IAM, which of the following predefined policy condition keys checks how long ago (in seconds) the MFA-validated security credentials making the request were issued using multi- factor authentication (MFA)?

- ☐ aws:MultiFactorAuthAge
- ☐ aws:MultiFactorAuthLast
- ☐ aws:MFAAge
- ☐ aws:MultiFactorAuthPrevious

Licensed to Berend Post berend.post@it-university.nl

Question 140

An organization has developed an application which provides a smarter shopping experience. They need to show a demonstration to various stakeholders who may not be able to access the in premise application so they decide to host a demo version of the application on AWS.
Consequently, they will need a fixed elastic IP attached automatically to the instance when it is launched.

In this scenario which of the below mentioned options will not help assign the elastic IP automatically?

- ☐ Write a script which will fetch the instance metadata on system boot and assign the public IP using that metadata.
- ☐ Provide an elastic IP in the user data and setup a bootstrapping script which will fetch that elastic IP and assign it to the instance.
- ☐ Create a controlling application which launches the instance and assigns the elastic IP based on the parameter provided when that instance is booted.
- ☐ Launch instance with VPC and assign an elastic IP to the primary network interface.

Question 141

Your company plans to host a large donation website on Amazon Web Services (AWS). You anticipate a large and undetermined amount of traffic that will create many database writes. To be certain that you do not drop any writes to a database hosted on AWS.

Which service should you use?

- ☐ Amazon RDS with provisioned IOPS up to the anticipated peak write throughput.
- ☐ Amazon Simple Queue Service (SQS) for capturing the writes and draining the queue to write to the database.
- ☐ Amazon ElastiCache to store the writes until the writes are committed to the database.
- ☐ Amazon DynamoDB with provisioned write throughput up to the anticipated peak write throughput.

Question 142

In AWS, which security aspects are the customer's responsibility? (Choose 4)

- ☐ Security Group and ACL (Access Control List) settings
- ☐ Decommissioning storage devices
- ☐ Patch management on the EC2 instance's operating system
- ☐ Life-cycle management of IAM credentials
- ☐ Controlling physical access to compute resources
- ☐ Encryption of EBS (Elastic Block Storage) volumes

Licensed to Berend Post berend.post@it-university.nl

Question 143

When using Numeric Conditions within IAM, short versions of the available comparators can be used instead of the more verbose versions.

Which of the following is the short version of the Numeric Condition "NumericLessThanEquals"?

- ☐ numlteq
- ☐ numlteql
- ☐ numltequals
- ☐ numeql

Question 144

As a part of building large applications in the AWS Cloud, the Solutions Architect is required to implement the perimeter security protection. Applications running on AWS have the following endpoints:

Application Load Balancer
Amazon API Gateway regional endpoint
Elastic IP address-based EC2 instances.
Amazon S3 hosted websites.
Classic Load Balancer

The Solutions Architect must design a solution to protect all of the listed web front ends and provide the following security capabilities:

DDoS protection
SQL injection protection
IP address whitelist/blacklist
HTTP flood protection
Bad bot scraper protection

How should the Solutions Architect design the solution?

- ☐ Deploy AWS WAF and AWS Shield Advanced on all web endpoints. Add AWS WAF rules to enforce the company's requirements.
- ☐ Deploy Amazon CloudFront in front of all the endpoints. The CloudFront distribution provides perimeter protection. Add AWS Lambda-based automation to provide additional security.
- ☐ Deploy Amazon CloudFront in front of all the endpoints. Deploy AWS WAF and AWS Shield Advanced. Add AWS WAF rules to enforce the company's requirements. Use AWS Lambda to automate and enhance the security posture.
- ☐ Secure the endpoints by using network ACLs and security groups and adding rules to enforce the company's requirements. Use AWS Lambda to automatically update the rules.

Licensed to Berend Post berend.post@it-university.nl

Question 145

You want to establish redundant VPN connections and customer gateways on your network by setting up a second VPN connection.

Which of the following will ensure that this functions correctly?

- ☐ The customer gateway IP address for the second VPN connection must be publicly accessible.
- ☐ The virtual gateway IP address for the second VPN connection must be publicly accessible.
- ☐ The customer gateway IP address for the second VPN connection must use dynamic routes.
- ☐ The customer gateway IP address for the second VPN connection must be privately accessible and be the same public IP address that you are using for the first VPN connection.

Question 146

You want to define permissions for a role in an IAM policy. Which of the following configuration formats should you use?

- ☐ An XML document written in the IAM Policy Language
- ☐ An XML document written in a language of your choice
- ☐ A JSON document written in the IAM Policy Language
- ☐ JSON document written in a language of your choice

Question 147

A Solutions Architect is working with a company that operates a standard three-tier web application in AWS. The web and application tiers run on Amazon EC2 and the database tier runs on Amazon RDS. The company is redesigning the web and application tiers to use Amazon API Gateway and AWS Lambda, and the company intends to deploy the new application within 6 months. The IT Manager has asked the Solutions Architect to reduce costs in the interim.

Which solution will be MOST cost effective while maintaining reliability?

- ☐ Use Spot Instances for the web tier, On-Demand Instances for the application tier, and Reserved Instances for the database tier.
- ☐ Use On-Demand Instances for the web and application tiers, and Reserved Instances for the database tier.
- ☐ Use Spot Instances for the web and application tiers, and Reserved Instances for the database tier.
- ☐ Use Reserved Instances for the web, application, and database tiers.

Licensed to Berend Post berend.post@it-university.nl

Question 148

An organization is trying to setup a VPC with Auto Scaling. Which configuration steps below is not required to setup AWS VPC with Auto Scaling?

- ☐ Configure the Auto Scaling group with the VPC ID in which instances will be launched.
- ☐ Configure the Auto Scaling Launch configuration with multiple subnets of the VPC to enable the Multi AZ feature.
- ☐ Configure the Auto Scaling Launch configuration which does not allow assigning a public IP to instances.
- ☐ Configure the Auto Scaling Launch configuration with the VPC security group.

Question 149

In the context of AWS IAM, identify a true statement about user passwords (login profiles).

- ☐ They must contain Unicode characters.
- ☐ They can contain any Basic Latin (ASCII) characters.
- ☐ They must begin and end with a forward slash (/).
- ☐ They cannot contain Basic Latin (ASCII) characters.

Question 150

In IAM, which of the following is true of temporary security credentials?

- ☐ Once you issue temporary security credentials, they cannot be revoked.
- ☐ None of these are correct.
- ☐ Once you issue temporary security credentials, they can be revoked only when the virtual MFA device is used.
- ☐ Once you issue temporary security credentials, they can be revoked.

Question 151

Which of the following components of AWS Data Pipeline specifies the business logic of your data management?

- ☐ Task Runner
- ☐ Pipeline definition
- ☐ AWS Direct Connect
- ☐ Amazon Simple Storage Service 9Amazon S3)

Licensed to Berend Post berend.post@it-university.nl

Question 152

Your customer wishes to deploy an enterprise application to AWS, which will consist of several web servers, several application servers and a small (50GB) Oracle database. Information is stored, both in the database and the file systems of the various servers. The backup system must support database recovery whole server and whole disk restores, and individual file restores with a recovery time of no more than two hours. They have chosen to use RDS Oracle as the database.

Which backup architecture will meet these requirements?

- ☐ Backup RDS using automated daily DB backups. Backup the EC2 instances using AMIs and supplement with file-level backup to S3 using traditional enterprise backup software to provide file level restore.
- ☐ Backup RDS using a Multi-AZ Deployment. Backup the EC2 instances using Amis, and supplement by copying file system data to S3 to provide file level restore.
- ☐ Backup RDS using automated daily DB backups. Backup the EC2 instances using EBS snapshots and supplement with file-level backups to Amazon Glacier using traditional enterprise backup software to provide file level restore.
- ☐ Backup RDS database to S3 using Oracle RMAN. Backup the EC2 instances using Amis, and supplement with EBS snapshots for individual volume restore.

Question 153

A company needs to cost-effectively persist small data records (up to 1 KiB) for up to 30 days. The data is read rarely. When reading the data, a 5-minute delay is acceptable.

Which of the following solutions achieve this goal? (Choose two.)

- ☐ Use Amazon S3 to collect multiple records in one S3 object. Use a lifecycle configuration to move data to Amazon Glacier immediately after write. Use expedited retrievals when reading the data.
- ☐ Write the records to Amazon Kinesis Data Firehose and configure Kinesis Data Firehose to deliver the data to Amazon S3 after 5 minutes. Set an expiration action at 30 days on the S3 bucket.
- ☐ Use an AWS Lambda function invoked via Amazon API Gateway to collect data for 5 minutes. Write data to Amazon S3 just before the Lambda execution stops.
- ☐ Write the records to Amazon DynamoDB configured with a Time To Live (TTL) of 30 days. Read data using the GetItem or BatchGetItem call.
- ☐ Write the records to an Amazon ElastiCache for Redis. Configure the Redis append-only file (AOF) persistence logs to write to Amazon S3. Recover from the log if the ElastiCache instance has failed.

Licensed to Berend Post berend.post@it-university.nl

Question 154

A user is planning to host a web server as well as an app server on a single EC2 instance which is a part of the public subnet of a VPC.

How can the user setup to have two separate public IPs and separate security groups for both the application as well as the web server?

- ☐ Launch VPC with two separate subnets and make the instance a part of both the subnets.
- ☐ Launch a VPC instance with two network interfaces. Assign a separate security group and elastic IP to them.
- ☐ Launch a VPC instance with two network interfaces. Assign a separate security group to each and AWS will assign a separate public IP to them.
- ☐ Launch a VPC with ELB such that it redirects requests to separate VPC instances of the public subnet.

Question 155

The CFO of a company wants to allow one of his employees to view only the AWS usage report page.

Which of the below mentioned IAM policy statements allows the user to have access to the AWS usage report page?

- ☐ "Effect": "Allow", "Action": ["Describe"], "Resource": "Billing"
- ☐ "Effect": "Allow", "Action": ["aws-portal: ViewBilling"], "Resource": "*"
- ☐ "Effect": "Allow", "Action": ["aws-portal: ViewUsage"], "Resource": "*"
- ☐ "Effect": "Allow", "Action": ["AccountUsage], "Resource": "*"

Question 156

True or false: In a CloudFormation template, you can reuse the same logical ID several times to reference the resources in other parts of the template.

- ☐ True, a logical ID can be used several times to reference the resources in other parts of the template.
- ☐ False, a logical ID must be unique within the template.
- ☐ False, you can mention a resource only once and you cannot reference it in other parts of a template.
- ☐ False, you cannot reference other parts of the template.

Licensed to Berend Post berend.post@it-university.nl

Question 157

Your company hosts a social media website for storing and sharing documents. The web application allows user to upload large files while resuming and pausing the upload as needed. Currently, files are uploaded to your PHP front end backed by Elastic Load Balancing and an autoscaling fleet of Amazon Elastic Compute Cloud (EC2) instances that scale upon average of bytes received (NetworkIn). After a file has been uploaded, it is copied to Amazon Simple Storage Service (S3). Amazon EC2 instances use an AWS Identity and Access Management (IAM) role that allows Amazon S3 uploads. Over the last six months, your user base and scale have increased significantly, forcing you to increase the Auto Scaling group's Max parameter a few times. Your CFO is concerned about rising costs and has asked you to adjust the architecture where needed to better optimize costs.

Which architecture change could you introduce to reduce costs and still keep your web application secure and scalable?

☐ Replace the Auto Scaling launch configuration to include c3.8xlarge instances; those instances can potentially yield a network throuthput of 10gbps.

☐ Re-architect your ingest pattern, have the app authenticate against your identity provider, and use your identity provider as a broker fetching temporary AWS credentials from AWS Secure Token Service (GetFederationToken). Securely pass the credentials and S3 endpoint/prefix to your app. Implement client-side logic to directly upload the file to Amazon S3 using the given credentials and S3 prefix.

☐ Re-architect your ingest pattern, and move your web application instances into a VPC public subnet. Attach a public IP address for each EC2 instance (using the Auto Scaling launch configuration settings). Use Amazon Route 53 Round Robin records set and HTTP health check to DNS load balance the app requests; this approach will significantly reduce the cost by bypassing Elastic Load Balancing.

☐ Re-architect your ingest pattern, have the app authenticate against your identity provider, and use your identity provider as a broker fetching temporary AWS credentials from AWS Secure Token Service (GetFederationToken). Securely pass the credentials and S3 endpoint/prefix to your app. Implement client-side logic that used the S3 multipart upload API to directly upload the file to Amazon S3 using the given credentials and S3 prefix.

Licensed to Berend Post berend.post@it-university.nl

Question 158

You have an application running on an EC2 instance which will allow users to download files from a private S3 bucket using a pre-signed URL. Before generating the URL, the application should verify the existence of the file in S3.

How should the application use AWS credentials to access the S3 bucket securely?

- ☐ Use the AWS account access keys; the application retrieves the credentials from the source code of the application.
- ☐ Create an IAM role for EC2 that allows list access to objects In the S3 bucket; launch the Instance with the role, and retrieve the role's credentials from the EC2 instance metadata.
- ☐ Create an IAM user for the application with permissions that allow list access to the S3 bucket; the application retrieves the 1AM user credentials from a temporary directory with permissions that allow read access only to the Application user.
- ☐ Create an IAM user for the application with permissions that allow list access to the S3 bucket; launch the instance as the IAM user, and retrieve the IAM user's credentials from the EC2 instance user data.

Question 159

You currently operate a web application. In the AWS US-East region. The application runs on an auto-scaled layer of EC2 instances and an RDS Multi-AZ database. Your IT security compliance officer has tasked you to develop a reliable and durable logging solution to track changes made to your EC2.IAM And RDS resources. The solution must ensure the integrity and confidentiality of your log data.

Which of these solutions would you recommend?

- ☐ Create a new CloudTrail trail with one new S3 bucket to store the logs and with the global services option selected. Use IAM roles S3 bucket policies and Multi Factor Authentication (MFA) Delete on the S3 bucket that stores your logs.
- ☐ Create a new CloudTrail with one new S3 bucket to store the logs Configure SNS to send log file delivery notifications to your management system. Use IAM roles and S3 bucket policies on the S3 bucket mat stores your logs.
- ☐ Create a new CloudTrail trail with an existing S3 bucket to store the logs and with the global services option selected. Use S3 ACLs and Multi Factor Authentication (MFA). Delete on the S3 bucket that stores your logs.
- ☐ Create three new CloudTrail trails with three new S3 buckets to store the logs one for the AWS Management console, one for AWS SDKs and one for command line tools. Use IAM roles and S3 bucket policies on the S3 buckets that store your logs.

Licensed to Berend Post berend.post@it-university.nl

Question 160

Regarding Amazon SNS, you can send notification messages to mobile devices through any of the following supported push notification services, EXCEPT:

- ☐ Microsoft Windows Mobile Messaging (MWMM)
- ☐ Google Cloud Messaging for Android (GCM)
- ☐ Amazon Device Messaging (ADM)
- ☐ Apple Push Notification Service (APNS)

Question 161

You have deployed a web application targeting a global audience across multiple AWS Regions under the domain name.example.com. You decide to use Route53 Latency-Based Routing to serve web requests to users from the region closest to the user. To provide business continuity in the event of server downtime you configure weighted record sets associated with two web servers in separate Availability Zones per region. Dunning a DR test you notice that when you disable all web servers in one of the regions Route53 does not automatically direct all users to the other region.

What could be happening? (Choose 2 answers)

- ☐ Latency resource record sets cannot be used in combination with weighted resource record sets.
- ☐ You did not setup an HTTP health check to one or more of the weighted resource record sets associated with me disabled web servers.
- ☐ The value of the weight associated with the latency alias resource record set in the region with the disabled servers is higher than the weight for the other region.
- ☐ One of the two working web servers in the other region did not pass its HTTP health check.
- ☐ You did not set "Evaluate Target Health" to "Yes" on the latency alias resource record set associated with example com in the region where you disabled the servers.

Licensed to Berend Post berend.post@it-university.nl

Question 162

A Solutions Architect has been asked to look at a company's Amazon Redshift cluster, which has quickly become an integral part of its technology and supports key business process. The Solutions Architect is to increase the reliability and availability of the cluster and provide options to ensure that if an issue arises, the cluster can either operate or be restored within four hours.

Which of the following solution options BEST addresses the business need in the most cost-effective manner?

- ☐ Ensure that the Amazon Redshift cluster has been set up to make use of Auto Scaling groups with the nodes in the cluster spread across multiple Availability Zones.
- ☐ Ensure that the Amazon Redshift cluster creation has been template using AWS CloudFormation so it can easily be launched in another Availability Zone and data populated from the automated Redshift back-ups stored in Amazon S3.
- ☐ Use Amazon Kinesis Data Firehose to collect the data ahead of ingestion into Amazon Redshift and create clusters using AWS CloudFormation in another region and stream the data to both clusters.
- ☐ Create two identical Amazon Redshift clusters in different regions (one as the primary, one as the secondary). Use Amazon S3 cross-region replication from the primary to secondary). Use Amazon S3 cross-region replication from the primary to secondary region, which triggers an AWS Lambda function to populate the cluster in the secondary region.

Question 163

You are tasked with moving a legacy application from a virtual machine running inside your datacenter to an Amazon VPC. Unfortunately, this app requires access to a number of on-premises services and no one who configured the app still works for your company. Even worse there's no documentation for it.

What will allow the application running inside the VPC to reach back and access its internal dependencies without being reconfigured? (Choose 3 answers)

- ☐ An AWS Direct Connect link between the VPC and the network housing the internal services.
- ☐ An Internet Gateway to allow a VPN connection.
- ☐ An Elastic IP address on the VPC instance
- ☐ An IP address space that does not conflict with the one on-premises
- ☐ Entries in Amazon Route 53 that allow the Instance to resolve its dependencies' IP addresses
- ☐ A VM Import of the current virtual machine

Licensed to Berend Post berend.post@it-university.nl

Question 164

Which status represents a failure state in AWS CloudFormation?

- ☐ ROLLBACK_IN_PROGRESS
- ☐ DELETE_IN_PROGRESS
- ☐ UPDATE_COMPLETE_CLEANUP_IN_PROGRESS
- ☐ REVIEW_IN_PROGRESS

Question 165

In Amazon RDS for PostgreSQL, you can provision up to 3TB storage and 30,000 IOPS per database instance. For a workload with 50% writes and 50% reads running on a cr1.8xlarge instance, you can realize over 25,000 IOPS for PostgreSQL. However, by provisioning more than this limit, you may be able to achieve:

- ☐ higher latency and lower throughput.
- ☐ lower latency and higher throughput.
- ☐ higher throughput only.
- ☐ higher latency only.

Question 166

When I/O performance is more important than fault tolerance, which of the following configurations should be used?

- ☐ SPAN 10
- ☐ RAID 1
- ☐ RAID 0
- ☐ NFS 1

Question 167

A customer has established an AWS Direct Connect connection to AWS. The link is up and routes are being advertised from the customer's end, however the customer is unable to connect from EC2 instances inside its VPC to servers residing in its datacenter.

Which of the following options provide a viable solution to remedy this situation? (Choose 2)

- ☐ Add a route to the route table with an iPsec VPN connection as the target.
- ☐ Enable route propagation to the virtual pinnate gateway (VGW).
- ☐ Enable route propagation to the customer gateway (CGW).
- ☐ Modify the route table of all Instances using the 'route' command.
- ☐ Modify the Instances VPC subnet route table by adding a route back to the customer's on-premises environment.

Licensed to Berend Post berend.post@it-university.nl

Question 168

Does Autoscaling automatically assign tags to resources?

- ☐ No, not unless they are configured via API.
- ☐ Yes, it does.
- ☐ Yes, by default.
- ☐ No, it does not.

Question 169

A 3-tier e-commerce web application is current deployed on-premises and will be migrated to AWS for greater scalability and elasticity. The web server currently shares read-only data using a network distributed file system. The app server tier uses a clustering mechanism for discovery and shared session state that depends on IP multicast. The database tier uses shared-storage clustering to provide database fall over capability, and uses several read slaves for scaling. Data on all servers and the distributed file system directory is backed up weekly to off-site tapes.

Which AWS storage and database architecture meets the requirements of the application?

- ☐ Web servers: store read-only data in S3, and copy from S3 to root volume at boot time. App servers: share state using a combination of DynamoDB and IP unicast. Database: use RDS with multi-AZ deployment and one or more read replicas. Backup: web servers, app servers, and database backed up weekly to Glacier using snapshots.
- ☐ Web servers: store read-only data in an EC2 NFS server; mount to each web server at boot time. App servers: share state using a combination of DynamoDB and IP multicast. Database: use RDS with multi-AZ deployment and one or more Read Replicas. Backup: web and app servers backed up weekly via AMIs, database backed up via DB snapshots.
- ☐ Web servers: store read-only data in S3, and copy from S3 to root volume at boot time. App servers: share state using a combination of DynamoDB and IP unicast. Database: use RDS with multi-AZ deployment and one or more Read Replicas. Backup: web and app servers backed up weekly via AMIs, database backed up via DB snapshots.
- ☐ Web servers: store read-only data in S3, and copy from S3 to root volume at boot time. App servers: share state using a combination of DynamoDB and IP unicast. Database: use RDS with multi-AZ deployment. Backup: web and app servers backed up weekly via AMIs, database backed up via DB snapshots.

Licensed to Berend Post berend.post@it-university.nl

Question 170

An organization is setting up a highly scalable application using Elastic Beanstalk. The organization is using ELB and RDS with VPC. The organization has public and private subnets within the cloud.

Which of the below mentioned configurations will not work in this scenario?

- ☐ To setup RDS in a private subnet and ELB in a public subnet.
- ☐ The configuration must have public and private subnets in the same AZ.
- ☐ The configuration must have two private subnets in separate AZs.
- ☐ The EC2 instance should have a public IP assigned to it.

Question 171

You have subscribed to the AWS Business and Enterprise support plan.
Your business has a backlog of problems, and you need about 20 of your IAM users to open technical support cases.

How many users can open technical support cases under the AWS Business and Enterprise support plan?

- ☐ 5 users
- ☐ 10 users
- ☐ Unlimited
- ☐ 1 user

Question 172

AWS CloudFormation _____ are special actions you use in your template to assign values to properties that are not available until runtime.

- ☐ intrinsic functions
- ☐ properties declarations
- ☐ output functions
- ☐ conditions declarations

Licensed to Berend Post berend.post@it-university.nl

Question 173

A large company experienced a drastic increase in its monthly AWS spend. This is after Developers accidentally launched Amazon EC2 instances in unexpected regions. The company has established practices around least privileges for Developers and controls access to on-premises resources using Active Directory groups. The company now want to control costs by restricting the level of access that Developers have to the AWS Management Console without impacting their productivity. The company would also like to allow Developers to launch Amazon EC2 in only one region, without limiting access to other services in any region.

How can this company achieve these new security requirements while minimizing the administrative burden on the Operations team?

- ☐ Set up SAML-based authentication tied to an IAM role that has an AdministrativeAccess managed policy attached to it. Attach a customer managed policy that denies access to Amazon EC2 in each region except for the one required.
- ☐ Create an IAM user for each Developer and add them to the developer IAM group that has the PowerUserAccess managed policy attached to it. Attach a customer managed policy that allows the Developers access to Amazon EC2 only in the required region.
- ☐ Set up SAML-based authentication tied to an IAM role that has a PowerUserAccess managed policy and a customer managed policy that deny all the Developers access to any AWS services except AWS Service Catalog. Within AWS Service Catalog, create a product containing only the EC2 resources in the approved region.
- ☐ Set up SAML-based authentication tied to an IAM role that has the PowerUserAccess managed policy attached to it. Attach a customer managed policy that denies access to Amazon EC2 in each region except for the one required.

Question 174

An organization is planning to host a Wordpress blog as well as joomla CMS on a single instance launched with VPC. The organization wants to create separate domains for each application using Route 53. The organization may have about ten instances each with these two applications. While launching each instance, the organization configured two separate network interfaces (primary + secondary ENI) with their own Elastic IPs to the instance. The suggestion was to use a public IP from AWS instead of an Elastic IP as the number of elastic IPs allocation per region is restricted in the account.

What action will you recommend to the organization?

- ☐ Only Elastic IP can be used by requesting limit increase, since AWS does not assign a public IP to an instance with multiple ENIs.
- ☐ AWS VPC does not attach a public IP to an ENI; so the only way is to use an Elastic IP.
- ☐ I agree with the suggestion but will prefer that the organization should use separate subnets with each ENI for different public IPs.
- ☐ I agree with the suggestion and it is recommended to use a public IP from AWS since the organization is going to use DNS with Route 53.

Licensed to Berend Post berend.post@it-university.nl

Question 175

A company has a serverless application comprised of Amazon CloudFront, Amazon API Gateway, and AWS Lambda functions. The current deployment process of the application code is to create a new version number of the Lambda function and run an AWS CLI script to update. If the new function version has errors, another CLI script reverts by deploying the previous working version of the function. The company would like to decrease the time to deploy new versions of the application logic provided by the Lambda functions, and also reduce the time to detect and revert when errors are identified.

How can this be accomplished?

☐ Create and deploy nested AWS CloudFormation stacks with the parent stack consisting of the AWS CloudFront distribution and API Gateway, and the child stack containing the Lambda function. For changes to Lambda, create an AWS CloudFormation change set and deploy; if errors are triggered, revert the AWS CloudFormation change set to the previous version.

☐ Use AWS SAM and built-in AWS CodeDeploy to deploy the new Lambda version, gradually shift traffic to the new version, and use pre-traffic and post-traffic test functions to verify code. Rollback if Amazon CloudWatch alarms are triggered.

☐ Refactor the AWS CLI scripts into a single script that deploys the new Lambda version. When deployment is completed, the script tests execute. If errors are detected, revert to the previous Lambda version.

☐ Create and deploy an AWS CloudFormation stack that consists of a new API Gateway endpoint that references the new Lambda version. Change the CloudFront origin to the new API Gateway endpoint, monitor errors and if detected, change the AWS CloudFront origin to the previous API Gateway endpoint.

Question 176

An administrator is using Amazon CloudFormation to deploy a three tier web application that consists of a web tier and application tier that will utilize Amazon DynamoDB for storage when creating the CloudFormation template.

Which of the following would allow the application instance access to the DynamoDB tables without exposing API credentials?

☐ Create an Identity and Access Management Role that has the required permissions to read and write from the required DynamoDB table and associate the Role to the application instances by referencing an instance profile.

☐ Use the Parameter section in the Cloud Formation template to nave the user input Access and Secret Keys from an already created IAM user that has me permissions required to read and write from the required DynamoDB table.

☐ Create an Identity and Access Management Role that has the required permissions to read and write from the required DynamoDB table and reference the Role in the instance profile property of the application instance.

☐ Create an identity and Access Management user in the CloudFormation template that has permissions to read and write from the required DynamoDB table, use the GetAtt

Licensed to Berend Post berend.post@it-university.nl

function to retrieve the Access and secret keys and pass them to the application instance through user-data.

Question 177

You are designing a social media site and are considering how to mitigate distributed denial-of-service (DDoS) attacks.

Which of the below are viable mitigation techniques? (Choose 3)

- ☐ Add multiple elastic network interfaces (ENIs) to each EC2 instance to increase the network bandwidth.
- ☐ Use dedicated instances to ensure that each instance has the maximum performance possible.
- ☐ Use an Amazon CloudFront distribution for both static and dynamic content.
- ☐ Use an Elastic Load Balancer with auto scaling groups at the web, app and Amazon Relational Database Service (RDS) tiers
- ☐ Add alert Amazon CloudWatch to look for high Network in and CPU utilization.
- ☐ Create processes and capabilities to quickly add and remove rules to the instance OS firewall.

Question 178

An organization, which has the AWS account ID as 999988887777, has created 50 IAM users. All the users are added to the same group ABC.
If the organization has enabled that each IAM user can login with the AWS console, which AWS login URL will the IAM users use??

- ☐ https://999988887777.aws.amazon.com/ABC/
- ☐ https://signin.aws.amazon.com/ABC/
- ☐ https://ABC.signin.aws.amazon.com/999988887777/console/
- ☐ https://999988887777.signin.aws.amazon.com/console/

Question 179

When using the AWS CLI for AWS CloudFormation, which of the following commands returns a description of the specified resource in the specified stack?

- ☐ describe-stack-events
- ☐ describe-stack-resource
- ☐ create-stack-resource
- ☐ describe-stack-returns

Licensed to Berend Post berend.post@it-university.nl

Question 180

A user has created a VPC with two subnets: one public and one private. The user is planning to run the patch update for the instances in the private subnet.

How can the instances in the private subnet connect to the internet?

- ☐ The private subnet can never connect to the internet
- ☐ Use NAT with an elastic IP
- ☐ Use the internet gateway with a private IP
- ☐ Allow outbound traffic in the security group for port 80 to allow internet updates

Question 181

An organization has 4 people in the IT operations team who are responsible to manage the AWS infrastructure. The organization wants to setup that each user will have access to launch and manage an instance in a zone which the other user cannot modify.

Which of the below mentioned options is the best solution to set this up?

- ☐ Create four AWS accounts and give each user access to a separate account.
- ☐ Create an IAM user and allow them permission to launch an instance of a different sizes only.
- ☐ Create four IAM users and four VPCs and allow each IAM user to have access to separate VPCs.
- ☐ Create a VPC with four subnets and allow access to each subnet for the individual IAM user.

Question 182

You are developing a new mobile application and are considering storing user preferences in AWS.2w This would provide a more uniform cross-device experience to users using multiple mobile devices to access the application. The preference data for each user is estimated to be 50KB in size Additionally 5 million customers are expected to use the application on a regular basis.

The solution needs to be cost-effective, highly available, scalable and secure, how would you design a solution to meet the above requirements?

- ☐ Setup an RDS MySQL instance in 2 availability zones to store the user preference data. Deploy a public facing application on a server in front of the database to manage security and access credentials
- ☐ Setup a DynamoDB table with an item for each user having the necessary attributes to hold the user preferences. The mobile application will query the user preferences directly from the DynamoDB table. Utilize STS. Web Identity Federation, and DynamoDB Fine Grained Access Control to authenticate and authorize access.
- ☐ Setup an RDS MySQL instance with multiple read replicas in 2 availability zones to store the user preference data .The mobile application will query the user preferences from

Licensed to Berend Post berend.post@it-university.nl

the read replicas. Leverage the MySQL user management and access privilege system to manage security and access credentials.

☐ Store the user preference data in S3 Setup a DynamoDB table with an item for each user and an item attribute pointing to the user' S3 object. The mobile application will retrieve the S3 URL from DynamoDB and then access the S3 object directly utilize STS, Web identity Federation, and S3 ACLs to authenticate and authorize access.

Question 183

After launching an instance that you intend to serve as a NAT (Network Address Translation) device in a public subnet you modify your route tables to have the NAT device be the target of internet bound traffic of your private subnet. When you try and make an outbound connection to the internet from an instance in the private subnet, you are not successful.

Which of the following steps could resolve the issue?

☐ Disabling the Source/Destination Check attribute on the NAT instance
☐ Attaching an Elastic IP address to the instance in the private subnet
☐ Attaching a second Elastic Network Interface (ENI) to the NAT instance, and placing it in the private subnet
☐ Attaching a second Elastic Network Interface (ENI) to the instance in the private subnet, and placing it in the public subnet

Question 184

Within the IAM service a GROUP is regarded as a:

☐ A collection of AWS accounts
☐ It's the group of EC2 machines that gain the permissions specified in the GROUP.
☐ There's no GROUP in IAM, but only USERS and RESOURCES.
☐ A collection of users.

Question 185

In Amazon Cognito, your mobile app authenticates with the Identity Provider (IdP) using the provider's SDK. Once the end user is authenticated with the IdP, the OAuth or OpenID Connect token returned from the IdP is passed by your app to Amazon Cognito, which returns a new _____ for the user and a set of temporary, limited-privilege AWS credentials.

☐ Cognito Key Pair
☐ Cognito API
☐ Cognito ID
☐ Cognito SDK

Licensed to Berend Post berend.post@it-university.nl

Question 186

While implementing the policy keys in AWS Direct Connect, if you use and the request comes from an Amazon EC2 instance, the instance's public IP address is evaluated to determine if access is allowed.

- ☐ aws:SecureTransport
- ☐ aws:EpochIP
- ☐ aws:SourceIp
- ☐ aws:CurrentTime

Question 187

Select the correct set of options. These are the initial settings for the default security group:

- ☐ Allow no inbound traffic, Allow all outbound traffic and Allow instances associated with this security group to talk to each other
- ☐ Allow all inbound traffic, Allow no outbound traffic and Allow instances associated with this security group to talk to each other
- ☐ Allow no inbound traffic, Allow all outbound traffic and Does NOT allow instances associated with this security group to talk to each other
- ☐ Allow all inbound traffic, Allow all outbound traffic and Does NOT allow instances associated with this security group to talk to each other

Question 188

After moving an E-Commerce website for a client from a dedicated server to AWS you have also set up auto scaling to perform health checks on the instances in your group and replace instances that fail these checks. Your client has come to you with his own health check system that he wants you to use as it has proved to be very useful prior to his site running on AWS.

What do you think would be an appropriate response to this given all that you know about auto scaling and CloudWatch?

- ☐ It is not possible to implement your own health check system due to compatibility issues.
- ☐ It is not possible to implement your own health check system. You need to use AWSs health check system.
- ☐ It is possible to implement your own health check system and then send the instance's health information directly from your system to CloudWatch but only in the US East (N. Virginia) region.
- ☐ It is possible to implement your own health check system and then send the instance's health information directly from your system to CloudWatch.

Licensed to Berend Post berend.post@it-university.nl

Question 189

Which of the following cannot be used to manage Amazon ElastiCache and perform administrative tasks?

- ☐ AWS software development kits (SDKs)
- ☐ Amazon S3
- ☐ ElastiCache command line interface (CLI)
- ☐ AWS CloudWatch

Question 190

A company manages more than 200 separate internet-facing web applications. All of the applications are deployed to AWS in a single AWS Region The fully qualified domain names (FQDNs) of all of the applications are made available through HTTPS using Application Load Balancers (ALBs). The ALBs are configured to use public SSL/TLS certificates.

A Solutions Architect needs to migrate the web applications to a multi-region architecture. All HTTPS services should continue to work without interruption.

Which approach meets these requirements?

- ☐ Request a certificate for each FQDN using AWS KMS. Associate the certificates with the ALBs in the primary AWS Region. Enable cross-region availability in AWS KMS for the certificates and associate the certificates with the ALBs in the secondary AWS Region.
- ☐ Generate the key pairs and certificate requests for each FQDN using AWS KMS. Associate the certificates with the ALBs in both the primary and secondary AWS Regions.
- ☐ Request a certificate for each FQDN using AWS Certificate Manager. Associate the certificates with the ALBs in both the primary and secondary AWS Regions.
- ☐ Request certificates for each FQDN in both the primary and secondary AWS Regions using AWS Certificate Manager. Associate the certificates with the corresponding ALBs in each AWS Region.

Licensed to Berend Post berend.post@it-university.nl

Question 191

A company is currently running a production workload on AWS that is very I/O intensive. Its workload consists of a single tier with 10 c4.8xlarge instances, each with 2 TB gp2 volumes. The number of processing jobs has recently increased, and latency has increased as well. The team realizes that they are constrained on the IOPS. For the application to perform efficiently, they need to increase the IOPS by 3,000 for each of the instances.

Which of the following designs will meet the performance goal MOST cost effectively?

- ☐ Change the type of Amazon EBS volume from gp2 to io1 and set provisioned IOPS to 9,000.
- ☐ Increase the size of the gp2 volumes in each instance to 3 TB.
- ☐ Create a new Amazon EFS file system and move all the data to this new file system. Mount this file system to all 10 instances.
- ☐ Create a new Amazon S3 bucket and move all the data to this new bucket. Allow each instance to access this S3 bucket and use it for storage.

Question 192

When configuring your customer gateway to connect to your VPC, the_____Association is established first between the virtual private gateway and customer gateway using the Pre-Shared Key as the authenticator.

- ☐ IPsec
- ☐ BGP
- ☐ IKE Security
- ☐ Tunnel

Question 193

Which of the following commands accepts binary data as parameters?

- ☐ --user-data
- ☐ –cipher text-key
- ☐ --aws-customer-key
- ☐ --describe-instances-user

Question 194

True or false: In CloudFormation, you cannot create an Amazon RDS DB instance from a snapshot.

- ☐ False, you can specify it in attributes
- ☐ False, you can specify it in condition
- ☐ False, you can specify it in resource properties
- ☐ True

Licensed to Berend Post berend.post@it-university.nl

Question 195

A company has an existing on-premises three-tier web application. The Linux web servers serve content from a centralized file share on a NAS server because the content is refreshed several times a day from various sources. The existing infrastructure is not optimized and the company would like to move to AWS in order to gain the ability to scale resources up and down in response to load. On-premises and AWS resources are connected using AWS Direct Connect.

How can the company migrate the web infrastructure to AWS without delaying the content refresh process?

- ☐ Create a cluster of web server Amazon EC2 instances behind a Classic Load Balancer on AWS. Share an Amazon EBS volume among all instances for the content. Schedule a periodic synchronization of this volume and the NAS server.
- ☐ Create an on-premises file gateway using AWS Storage Gateway to replace the NAS server and replicate content to AWS. On the AWS side, mount the same Storage Gateway bucket to each web server Amazon EC2 instance to serve the content.
- ☐ Expose an Amazon EFS share to on-premises users to serve as the NAS serve. Mount the same EFS share to the web server Amazon EC2 instances to serve the content.
- ☐ Create web server Amazon EC2 instances on AWS in an Auto Scaling group. Configure a nightly process where the web server instances are updated from the NAS server.

Question 196

A user has created a launch configuration for Auto Scaling where CloudWatch detailed monitoring is disabled. The user wants to now enable detailed monitoring.

How can the user achieve this?

- ☐ Update the Launch config with CLI to set InstanceMonitoringDisabled = false
- ☐ The user should change the Auto Scaling group from the AWS console to enable detailed monitoring
- ☐ Create a new Launch Config with detail monitoring enabled and update the Auto Scaling group
- ☐ Update the Launch config with CLI to set InstanceMonitoring.Enabled = true

Licensed to Berend Post berend.post@it-university.nl

Question 197

An organization is having a VPC for the HR department, and another VPC for the Admin department. The HR department requires access to all the instances running in the Admin VPC while the Admin department requires access to all the resources in the HR department.

How can the organization setup this scenario?

- ☐ Setup VPC peering between the VPCs of Admin and HR.
- ☐ Setup ACL with both VPCs which will allow traffic from the CIDR of the other VPC.
- ☐ Setup the security group with each VPC which allows traffic from the CIDR of another VPC.
- ☐ It is not possible to connect resources of one VPC from another VPC.

Question 198

An e-commerce company is revamping its IT infrastructure and is planning to use AWS services. The company's CIO has asked a Solutions Architect to design a simple, highly available, and loosely coupled order processing application. The application is responsible for receiving and processing orders before storing them in an Amazon DynamoDB table. The application has a sporadic traffic pattern and should be able to scale during marketing campaigns to process the orders with minimal delays.

Which of the following is the MOST reliable approach to meet the requirements?

- ☐ Receive the orders in an Amazon EC2-hosted database and use EC2 instances to process them.
- ☐ Receive the orders in an Amazon SQS queue and trigger an AWS Lambda function to process them.
- ☐ Receive the orders using the AWS Step Functions program and trigger an Amazon ECS container to process them.
- ☐ Receive the orders in Amazon Kinesis Data Streams and use Amazon EC2 instances to process them.

Question 199

How can you check the operational validity of your AWS CloudFormation template?

- ☐ To check the operational validity, you need to attempt to create the stack.
- ☐ There is no way to check the operational validity of your AWS CloudFormation template.
- ☐ To check the operational validity, you need a sandbox or test area for AWS CloudFormation stacks.
- ☐ To check the operational validity, you need to use the aws cloudformation validate-template command.

Licensed to Berend Post berend.post@it-university.nl

Question 200

A read only news reporting site with a combined web and application tier and a database tier that receives large and unpredictable traffic demands must be able to respond to these traffic fluctuations automatically.

What AWS services should be used meet these requirements?

- ☐ Stateless instances for the web and application tier synchronized using ElastiCache Memcached in an autoscaimg group monitored with CloudWatch and RDS with read replicas.
- ☐ Stateful instances for the web and application tier in an autoscaling group monitored with CloudWatch and RDS with read replicas.
- ☐ Stateful instances for the web and application tier in an autoscaling group monitored with CloudWatch and multi-AZ RDS.
- ☐ Stateless instances for the web and application tier synchronized using ElastiCache Memcached in an autoscaling group monitored with CloudWatch and multi-AZ RDS.

Question 201

A large real-estate brokerage is exploring the option of adding a cost-effective location based alert to their existing mobile application. The application backend infrastructure currently runs on AWS. Users who opt in to this service will receive alerts on their mobile device regarding real-estate otters in proximity to their location. For the alerts to be relevant delivery time needs to be in the low minute count the existing mobile app has 5 million users across the US.

Which one of the following architectural suggestions would you make to the customer?

- ☐ The mobile application will submit its location to a web service endpoint utilizing Elastic Load Balancing and EC2 instances; DynamoDB will be used to store and retrieve relevant offers EC2 instances will communicate with mobile earners/device providers to push alerts back to mobile application.
- ☐ Use AWS DirectConnect or VPN to establish connectivity with mobile carriers EC2 instances will receive the mobile applications location through carrier connection: RDS will be used to store and relevant offers. EC2 instances will communicate with mobile carriers to push alerts back to the mobile application.
- ☐ The mobile application will send device location using SQS. EC2 instances will retrieve the relevant others from DynamoDB. AWS Mobile Push will be used to send offers to the mobile application.
- ☐ The mobile application will send device location using AWS Mobile Push EC2 instances will retrieve the relevant offers from DynamoDB. EC2 instances will communicate with mobile carriers/device providers to push alerts back to the mobile application.

Licensed to Berend Post berend.post@it-university.nl

Question 202

Your startup wants to implement an order fulfillment process for selling a personalized gadget that needs an average of 3-4 days to produce with some orders taking up to 6 months you expect 10 orders per day on your first day. 1000 orders per day after 6 months and 10,000 orders after 12 months.
Orders coming in are checked for consistency men dispatched to your manufacturing plant for production quality control packaging shipment and payment processing If the product does not meet the quality standards at any stage of the process employees may force the process to repeat a step Customers are notified via email about order status and any critical issues with their orders such as payment failure.
Your base architecture includes AWS Elastic Beanstalk for your website with an RDS MySQL instance for customer data and orders.

How can you implement the order fulfillment process while making sure that the emails are delivered reliably?

☐ Add a business process management application to your Elastic Beanstalk app servers and re-use the ROS database for tracking order status use one of the Elastic Beanstalk instances to send emails to customers.
☐ Use SWF with an Auto Scaling group of activity workers and a decider instance in another Auto Scaling group with min/max=1 Use the decider instance to send emails to customers.
☐ Use SWF with an Auto Scaling group of activity workers and a decider instance in another Auto Scaling group with min/max=1 use SES to send emails to customers.
☐ Use an SQS queue to manage all process tasks Use an Auto Scaling group of EC2 Instances that poll the tasks and execute them. Use SES to send emails to customers.

Question 203

A company had a tight deadline to migrate its on-premises environment to AWS. It moved over Microsoft SQL Servers and Microsoft Windows Servers using the virtual machine import/export service and rebuild other applications native to the cloud. The team created both Amazon EC2 databases and used Amazon RDS. Each team in the company was responsible for migrating their applications, and would like suggestions on reducing its AWS spend.

Which steps should a Solutions Architect take to reduce costs?

☐ Enable AWS Business Support and review AWS Trusted Advisor's cost checks. Create Amazon EC2 Auto Scaling groups for applications that experience fluctuating demand. Save AWS Simple Monthly Calculator reports in Amazon S3 for trend analysis. Create a master account under Organizations and have teams join for consolidating billing.
☐ Enable Cost Explorer and AWS Business Support Reserve Amazon EC2 and Amazon RDS DB instances. Use Amazon CloudWatch and AWS Trusted Advisor for monitoring and to receive cost-savings suggestions. Create a master account under Organizations and have teams join for consolidated billing.
☐ Create an AWS Lambda function that changes the instance size based on Amazon CloudWatch alarms. Reserve instances based on AWS Simple Monthly Calculator

Licensed to Berend Post berend.post@it-university.nl

suggestions. Have an AWS Well-Architected framework review and apply recommendations. Create a master account under Organizations and have teams join for consolidated billing.

☐ Create a budget and monitor for costs exceeding the budget. Create Amazon EC2 Auto Scaling groups for applications that experience fluctuating demand. Create an AWS Lambda function that changes instance sizes based on Amazon CloudWatch alarms. Have each team upload their bill to an Amazon S3 bucket for analysis of team spending. Use Spot instances on nightly batch processing jobs.

Licensed to Berend Post berend.post@it-university.nl

Question 204

A company has created an account for individual Development teams, resulting in a total of 200 accounts. All accounts have a single virtual private cloud (VPC) in a single region with multiple microservices running in Docker containers that need to communicate with microservices in other accounts. The Security team requirements state that these microservices must not traverse the public internet, and only certain internal services should be allowed to call other individual services. If there is any denied network traffic for a service, the Security team must be notified of any denied requests, including the source IP.

How can connectivity be established between service while meeting the security requirements?

☐ Create a VPC peering connection between the VPCs. Use security groups on the instances to allow traffic from the security group IDs that are permitted to call the microservice. Apply network ACLs to and allow traffic from the local VPC and peered VPCs only. Within the task definition in Amazon ECS for each of the microservices, specify a log configuration by using the awslogs driver. Within Amazon CloudWatch Logs, create a metric filter and alarm off of the number of HTTP 403 responses. Create an alarm when the number of messages exceeds a threshold set by the Security team.

☐ Ensure that no CIDR ranges are overlapping, and attach a virtual private gateway (VGW) to each VPC. Provision an IPsec tunnel between each VGW and enable route propagation on the route table. Configure security groups on each service to allow the CIDR ranges of the VPCs on the other accounts. Enable VPC Flow Logs, and use an Amazon CloudWatch Logs subscription filter for rejected traffic. Create an IAM role and allow the Security team to call the AssumeRole action for each account.

☐ Deploy a transit VPC by using third-party marketplace VPN appliances running on Amazon EC2, dynamically routed VPN connections between the VPN appliance, and the virtual private gateways (VGWs) attached to each VPC within the region. Adjust network ACLs to allow traffic from the local VPC only. Apply security groups to the microservices to allow traffic from the VPN appliances only. Install the awslogs agent on each VPN appliance, and configure logs to forward to Amazon CloudWatch Logs in the security account for the Security team to access.

☐ Create a Network Load Balancer (NLB) for each microservice. Attach the NLB to a PrivateLink endpoint service and whitelist the accounts that will be consuming this service. Create an interface endpoint in the consumer VPC and associate a security group that allows only the security group IDs of the services authorized to call the producer service. On the producer services, create security groups for each microservice and allow only the CIDR range the allowed services. Create VPC Flow Logs on each VPC to capture rejected traffic that will be delivered to an Amazon CloudWatch Logs group. Create a CloudWatch Logs subscription that streams the log data to a security account.

Licensed to Berend Post berend.post@it-university.nl

Question 205

If a single condition within an IAM policy includes multiple values for one key, it will be evaluated using a logical_____.

- ☐ OR
- ☐ NAND
- ☐ NOR
- ☐ AND

Question 206

For AWS CloudFormation, which stack state refuses UpdateStack calls?

- ☐ UPDATE_ROLLBACK_FAILED
- ☐ UPDATE_ROLLBACK_COMPLETE
- ☐ UPDATE_COMPLETE
- ☐ CREATE_COMPLETE

Question 207

A user has created a VPC with public and private subnets. The VPC has CIDR 20.0.0.0/16. The private subnet uses CIDR 20.0.1.0/24 and the public subnet uses CIDR 20.0.0.0/24. The user is planning to host a web server in the public subnet (port 80) and a DB server in the private subnet (port 3306). The user is configuring a security group of the NAT instance.

Which of the below mentioned entries is not required in NAT's security group for the database servers to connect to the Internet for software updates?

- ☐ For Outbound allow Destination: 0.0.0.0/0 on port 443
- ☐ For Inbound allow Source: 20.0.1.0/24 on port 80
- ☐ For Inbound allow Source: 20.0.0.0/24 on port 80
- ☐ For Outbound allow Destination: 0.0.0.0/0 on port 80

Licensed to Berend Post berend.post@it-university.nl

Question 208

An organization is setting up an application on AWS to have both High Availability (HA) and Disaster Recovery (DR). The organization wants to have both Recovery point objective (RPO) and Recovery time objective (RTO) of 10 minutes.

Which of the below mentioned service configurations does not help the organization achieve the said RPO and RTO?

- ☐ Take a snapshot of the data every 10 minutes and copy it to the other region.
- ☐ Use an elastic IP to assign to a running instance and use Route 53 to map the user's domain with that IP.
- ☐ Create ELB with multi-region routing to allow automated failover when required.
- ☐ Use an AMI copy to keep the AMI available in other regions.

Question 209

Your application provides data transformation services. Files containing data to be transformed are first uploaded to Amazon S3 and then transformed by a fleet of spot EC2 instances. Files submitted by your premium customers must be transformed with the highest priority.

How should you implement such a system?

- ☐ Use a DynamoDB table with an attribute defining the priority level. Transformation instances will scan the table for tasks, sorting the results by priority level.
- ☐ Use Route 53 latency based-routing to send high priority tasks to the closest transformation instances.
- ☐ Use two SQS queues, one for high priority messages, the other for default priority. Transformation instances first poll the high priority queue; if there is no message, they poll the default priority queue.
- ☐ Use a single SQS queue. Each message contains the priority level. Transformation instances poll high-priority messages first.

Question 210

What happens when Dedicated instances are launched into a VPC?

- ☐ If you launch an instance into a VPC that has an instance tenancy of dedicated, you must manually create a Dedicated instance.
- ☐ If you launch an instance into a VPC that has an instance tenancy of dedicated, your instance is created as a Dedicated instance, only based on the tenancy of the instance.
- ☐ If you launch an instance into a VPC that has an instance tenancy of dedicated, your instance is automatically a Dedicated instance, regardless of the tenancy of the instance.
- ☐ None of these are true.

Licensed to Berend Post berend.post@it-university.nl

Question 211

An organization is having an application which can start and stop an EC2 instance as per schedule. The organization needs the MAC address of the instance to be registered with its software. The instance is launched in EC2-CLASSIC.

How can the organization update the MAC registration every time an instance is booted?

- ☐ The organization should write a boot strapping script which will get the MAC address from the instance metadata and use that script to register with the application.
- ☐ The organization should provide a MAC address as a part of the user data. Thus, whenever the instance is booted the script assigns the fixed MAC address to that instance.
- ☐ The instance MAC address never changes. Thus, it is not required to register the MAC address every time.
- ☐ AWS never provides a MAC address to an instance; instead the instance ID is used for identifying the instance for any software registration.

Question 212

You're running an application on-premises due to its dependency on non-x86 hardware and want to use AWS for data backup. Your backup application is only able to write to POSIX-compatible block-based storage. You have 140TB of data and would like to mount it as a single folder on your file server. Users must be able to access portions of this data while the backups are taking place.

What backup solution would be most appropriate for this use case?

- ☐ Use Storage Gateway and configure it to use Gateway Cached volumes.
- ☐ Configure your backup software to use S3 as the target for your data backups.
- ☐ Configure your backup software to use Glacier as the target for your data backups.
- ☐ Use Storage Gateway and configure it to use Gateway Stored volumes.

Question 213

Is there any way to own a direct connection to Amazon Web Services?

- ☐ No, AWS only allows access from the public Internet.
- ☐ No, you can create an encrypted tunnel to VPC, but you cannot own the connection.
- ☐ Yes, you can via Amazon Dedicated Connection
- ☐ Yes, you can via AWS Direct Connect.

Licensed to Berend Post berend.post@it-university.nl

Question 214

Your company has recently extended its datacenter into a VPC on AWS to add burst computing capacity as needed Members of your Network Operations Center need to be able to go to the AWS Management Console and administer Amazon EC2 instances as necessary. You don't want to create new IAM users for each NOC member and make those users sign in again to the AWS Management Console.

Which option below will meet the needs for your NOC members?

- ☐ Use OAuth 2.0 to retrieve temporary AWS security credentials to enable your NOC members to sign in to the AWS Management Console.
- ☐ Use web Identity Federation to retrieve AWS temporary security credentials to enable your NOC members to sign in to the AWS Management Console.
- ☐ Use your on-premises SAML 2.0-compliant identity provider (IDP) to grant the NOC members federated access to the AWS Management Console via the AWS single sign-on (SSO) endpoint.
- ☐ Use your on-premises SAML2.0-compliam identity provider (IDP) to retrieve temporary security credentials to enable NOC members to sign in to the AWS Management Console.

Question 215

A company is migrating a subset of its application APIs from Amazon EC2 instances to run on a serverless infrastructure. The company has set up Amazon API Gateway, AWS Lambda, and Amazon DynamoDB for the new application. The primary responsibility of the Lambda function is to obtain data from a third-party Software as a Service (SaaS) provider. For consistency, the Lambda function is attached to the same virtual private cloud (VPC) as the original EC2 instances.

Test users report an inability to use this newly moved functionality, and the company is receiving 5xx errors from API Gateway. Monitoring reports from the SaaS provider shows that the requests never made it to its systems. The company notices that Amazon CloudWatch Logs are being generated by the Lambda functions. When the same functionality is tested against the EC2 systems, it works as expected.

What is causing the issue?

- ☐ Lambda is in a subnet that does not have a NAT gateway attached to it to connect to the SaaS provider.
- ☐ The end-user application is misconfigured to continue using the endpoint backed by EC2 instances.
- ☐ The throttle limit set on API Gateway is too low and the requests are not making their way through.
- ☐ API Gateway does not have the necessary permissions to invoke Lambda.

Licensed to Berend Post berend.post@it-university.nl

Question 216

You have a periodic image analysis application that gets some files in input, analyzes them and tor each file writes some data in output to a ten file the number of files in input per day is high and concentrated in a few hours of the day.
Currently you have a server on EC2 with a large EBS volume that hosts the input data and the results. It takes almost 20 hours per day to complete the process.

What services could be used to reduce the elaboration time and improve the availability of the solution?

- ☐ S3 to store I/O files. SQS to distribute elaboration commands to a group of hosts working in parallel. Auto scaling to dynamically size the group of hosts depending on the length of the SQS queue
- ☐ EBS with Provisioned IOPS (PIOPS) to store I/O files. SNS to distribute elaboration commands to a group of hosts working in parallel Auto Scaling to dynamically size the group of hosts depending on the number of SNS notifications
- ☐ S3 to store I/O files, SNS to distribute evaporation commands to a group of hosts working in parallel. Auto scaling to dynamically size the group of hosts depending on the number of SNS notifications
- ☐ EBS with Provisioned IOPS (PIOPS) to store I/O files SQS to distribute elaboration commands to a group of hosts working in parallel Auto Scaling to dynamically size the group ot hosts depending on the length of the SQS queue.

Question 217

A user has created a VPC with a public subnet. The user has terminated all the instances which are part of the subnet.

Which of the below mentioned statements is true with respect to this scenario?

- ☐ The subnet to which the instances were launched with will be deleted
- ☐ When the user launches a new instance it cannot use the same subnet
- ☐ The user cannot delete the VPC since the subnet is not deleted
- ☐ Secondary network interfaces attached to the terminated instances may persist.

Licensed to Berend Post berend.post@it-university.nl

Question 218

A government client needs you to set up secure cryptographic key storage for some of their extremely confidential data. You decide that the AWS CloudHSM is the best service for this. However, there seem to be a few pre-requisites before this can happen, one of those being a security group that has certain ports open.

Which of the following is correct in regards to those security groups?

- ☐ A security group that has no ports open to your network.
- ☐ A security group that has only port 3389 (for RDP) open to your network.
- ☐ A security group that has only port 22 (for SSH) open to your network.
- ☐ A security group that has port 22 (for SSH) or port 3389 (for RDP) open to your network.

Question 219

You have been given the task to define multiple AWS Data Pipeline schedules for different activities in the same pipeline.

Which of the following would successfully accomplish this task?

- ☐ Creating multiple pipeline definition files
- ☐ Defining multiple pipeline definitions in your schedule objects file and associating the desired schedule to the correct activity via its schedule field
- ☐ Defining multiple schedule objects in your pipeline definition file and associating the desired schedule to the correct activity via its schedule field
- ☐ Defining multiple schedule objects in the schedule field

Question 220

In a VPC, can you modify a set of DHCP options after you create them?

- ☐ Yes, you can modify a set of DHCP options within 48 hours after creation and there are no VPCs associated with them.
- ☐ Yes, you can modify a set of DHCP options any time after you create them.
- ☐ No, you can't modify a set of DHCP options after you create them.
- ☐ Yes, you can modify a set of DHCP options within 24 hours after creation.

Licensed to Berend Post berend.post@it-university.nl

Question 221

A Solutions Architect is designing a multi-account structure that has 10 existing accounts. The design must meet the following requirements:

Consolidate all accounts into one organization.
Allow full access to the Amazon EC2 service from the master account and the secondary accounts.
Minimize the effort required to add additional secondary accounts.

Which combination of steps should be included in the solution? (Choose two.)

- ☐ Create an organization from the master account. Send invitations to the secondary accounts from the master account. Accept the invitations and create an OU.
- ☐ Create an organization from the master account. Send a join request to the master account from each secondary account. Accept the requests and create an OU.
- ☐ Create a VPC peering connection between the master account and the secondary accounts. Accept the request for the VPC peering connection.
- ☐ Create a service control policy (SCP) that enables full EC2 access, and attach the policy to the OU.
- ☐ Create a full EC2 access policy and map the policy to a role in each account. Trust every other account to assume the role.

Question 222

An organization is undergoing a security audit. The auditor wants to view the AWS VPC configurations as the organization has hosted all the applications in the AWS VPC. The auditor is from a remote place and wants to have access to AWS to view all the VPC records.

How can the organization meet the expectations of the auditor without compromising on the security of their AWS infrastructure?

- ☐ The organization should not accept the request as sharing the credentials means compromising on security.
- ☐ Create an IAM role which will have read only access to all EC2 services including VPC and assign that role to the auditor.
- ☐ Create an IAM user who will have read only access to the AWS VPC and share those credentials with the auditor.
- ☐ The organization should create an IAM user with VPC full access but set a condition that will not allow to modify anything if the request is from any IP other than the organization's data center.

Licensed to Berend Post berend.post@it-university.nl

Question 223

Over which of the following Ethernet standards does AWS Direct Connect link your internal network to an AWS Direct Connect location?

- ☐ Single mode fiber-optic cable
- ☐ Multi-mode fiber-optic cable
- ☐ Shielded balanced copper cable
- ☐ Twisted pair cable

Question 224

How can an EBS volume that is currently attached to an EC2 instance be migrated from one Availability Zone to another?

- ☐ Detach the volume and attach it to another EC2 instance in the other AZ.
- ☐ Simply create a new volume in the other AZ and specify the original volume as the source.
- ☐ Create a snapshot of the volume, and create a new volume from the snapshot in the other AZ.
- ☐ Detach the volume, then use the ec2-migrate-volume command to move it to another AZ.

Question 225

A customer has a website which shows all the deals available across the market. The site experiences a load of 5 large EC2 instances generally.
However, a week before Thanksgiving vacation they encounter a load of almost 20 large instances. The load during that period varies over the day based on the office timings.

Which of the below mentioned solutions is cost effective as well as help the website achieve better performance?

- ☐ Setup to run 10 instances during the pre-vacation period and only scale up during the office time by launching 10 more instances using the AutoScaling schedule.
- ☐ Keep only 10 instances running and manually launch 10 instances every day during office hours.
- ☐ During the pre-vacation period setup 20 instances to run continuously.
- ☐ During the pre-vacation period setup a scenario where the organization has 15 instances running and 5 instances to scale up and down using Auto Scaling based on the network I/O policy.

Licensed to Berend Post berend.post@it-university.nl

Question 226

A company has an application that runs a web service on Amazon EC2 instances and stores .jpg images in Amazon S3. The web traffic has a predictable baseline, but often demand spikes unpredictably for short periods of time. The application is loosely coupled and stateless. The .jpg images stored in Amazon S3 are accessed frequently for the first 15 to 20 days, they are seldom accessed thereafter but always need to be immediately available. The CIO has asked to find ways to reduce costs.

Which of the following options will reduce costs? (Choose two.)

- ☐ Purchase Reserved instances for baseline capacity requirements and use On-Demand instances for the demand spikes.
- ☐ Configure a lifecycle policy to move the .jpg images on Amazon S3 to S3 IA after 30 days.
- ☐ Use On-Demand instances for baseline capacity requirements and use Spot Fleet instances for the demand spikes.
- ☐ Configure a lifecycle policy to move the .jpg images on Amazon S3 to Amazon Glacier after 30 days.
- ☐ Create a script that checks the load on all web servers and terminates unnecessary On-Demand instances.

Question 227

A user wants to create a public subnet in VPC and launch an EC2 instance within it. The user has not selected the option to assign a public IP address while launching the instance.

Which of the below mentioned statements is true with respect to this scenario?

- ☐ The instance will always have a public DNS attached to the instance by default
- ☐ The user would need to create a default route to IGW in subnet's route table and then attach an elastic IP to the instance to connect from the internet
- ☐ The user can directly attach an elastic IP to the instance
- ☐ The instance will never launch if the public IP is not assigned

Licensed to Berend Post berend.post@it-university.nl

Question 228

You are running a news website in the eu-west-1 region that updates every 15 minutes. The website has a world-wide audience. It uses an Auto Scaling group behind an Elastic Load Balancer and an Amazon RDS database. Static content resides on Amazon S3, and is distributed through Amazon CloudFront. Your Auto Scaling group is set to trigger a scale up event at 60% CPU utilization. You use an Amazon RDS extra large DB instance with 10.000 Provisioned IOPS, its CPU utilization is around 80%, while freeable memory is in the 2 GB range.

Web analytics reports show that the average load time of your web pages is around 1.5 to 2 seconds, but your SEO consultant wants to bring down the average load time to under 0.5 seconds.

How would you improve page load times for your users? (Choose 3 answers)

- ☐ Lower the scale up trigger of your Auto Scaling group to 30% so it scales more aggressively.
- ☐ Add an Amazon ElastiCache caching layer to your application for storing sessions and frequent DB queries
- ☐ Configure Amazon CloudFront dynamic content support to enable caching of re-usable content from your site
- ☐ Switch the Amazon RDS database to the high memory extra large Instance type
- ☐ Set up a second installation in another region, and use the Amazon Route 53 latency-based routing feature to select the right region.

Question 229

A user wants to configure AutoScaling which scales up when the CPU utilization is above 70% and scales down when the CPU utilization is below 30%.

How can the user configure AutoScaling for the above mentioned condition?

- ☐ Configure ELB to notify AutoScaling on load increase or decrease
- ☐ Use AutoScaling with a schedule
- ☐ Use AutoScaling by manually modifying the desired capacity during a condition
- ☐ Use dynamic AutoScaling with a policy

Licensed to Berend Post berend.post@it-university.nl

Question 230

A company is using an Amazon CloudFront distribution to distribute both static and dynamic content from a web application running behind an Application Load Balancer. The web application requires user authorization and session tracking for dynamic content. The CloudFront distribution has a single cache behavior configured to forward the Authorization, Host, and User-Agent HTTP whitelist headers and a session cookie to the origin. All other cache behavior settings are set to their default value.

A valid ACM certificate is applied to the CloudFront distribution with a matching CNAME in the distribution settings. The ACM certificate is also applied to the HTTPS listener for the Application Load Balancer. The CloudFront origin protocol policy is set to HTTPS only. Analysis of the cache statistics report shows that the miss rate for this distribution is very high.

What can the Solutions Architect do to improve the cache hit rate for this distribution without causing the SSL/TLS handshake between CloudFront and the Application Load Balancer to fail?

☐ Create two cache behaviors for static and dynamic content. Remove the User-Agent and Host HTTP headers from the whitelist headers section on both if the cache behaviors. Remove the session cookie from the whitelist cookies section and the Authorization HTTP header from the whitelist headers section for cache behavior configured for static content.
☐ Remove the User-Agent and Authorization HTTPS headers from the whitelist headers section of the cache behavior. Then update the cache behavior to use presigned cookies for authorization.
☐ Remove the Host HTTP header from the whitelist headers section and remove the session cookie from the whitelist cookies section for the default cache behavior. Enable automatic object compression and use Lambda@Edge viewer request events for user authorization.
☐ Create two cache behaviors for static and dynamic content. Remove the User-Agent HTTP header from the whitelist headers section on both of the cache behaviors. Remove the session cookie from the whitelist cookies section and the Authorization HTTP header from the whitelist headers section for cache behavior configured for static content.

Question 231

A bucket owner has allowed another account's IAM users to upload or access objects in his bucket. The IAM user of Account A is trying to access an object created by the IAM user of account B. What will happen in this scenario?

☐ It is not possible to give permission to multiple IAM users
☐ AWS S3 will verify proper rights given by the owner of Account A, the bucket owner as well as by the IAM user B to the object
☐ The bucket policy may not be created as S3 will give error due to conflict of Access Rights
☐ It is not possible that the IAM user of one account accesses objects of the other IAM user

Licensed to Berend Post berend.post@it-university.nl

Question 232

A benefits enrollment company is hosting a 3-tier web application running in a VPC on AWS which includes a NAT (Network Address Translation) instance in the public Web tier. There is enough provisioned capacity for the expected workload tor the new fiscal year benefit enrollment period plus some extra overhead Enrollment proceeds nicely for two days and then the web tier becomes unresponsive, upon investigation using CloudWatch and other monitoring tools it is discovered that there is an extremely large and unanticipated amount of inbound traffic coming from a set of 15 specific IP addresses over port 80 from a country where the benefits company has no customers. The web tier instances are so overloaded that benefit enrollment administrators cannot even SSH into them.

Which activity would be useful in defending against this attack?

- ☐ Create a custom route table associated with the web tier and block the attacking IP addresses from the IGW (Internet Gateway)
- ☐ Change the EIP (Elastic IP Address) of the NAT instance in the web tier subnet and update the Main Route Table with the new EIP
- ☐ Create 15 Security Group rules to block the attacking IP addresses over port 80
- ☐ Create an inbound NACL (Network Access control list) associated with the web tier subnet with deny rules to block the attacking IP addresses

Question 233

To scale out the AWS resources using manual AutoScaling, which of the below mentioned parameters should the user change?

- ☐ Current capacity
- ☐ Desired capacity
- ☐ Preferred capacity
- ☐ Maximum capacity

Question 234

A user is trying to create a vault in AWS Glacier. The user wants to enable notifications.

In which of the below mentioned options can the user enable the notifications from the AWS console?

- ☐ Glacier does not support the AWS console
- ☐ Archival Upload Complete
- ☐ Vault Upload Job Complete
- ☐ Vault Inventory Retrieval Job Complete

Licensed to Berend Post berend.post@it-university.nl

Question 235

What is a circular dependency in AWS CloudFormation?

- ☐ When Nested Stacks depend on each other.
- ☐ When Resources form a Depend On loop.
- ☐ When a Template references an earlier version of itself.
- ☐ When a Template references a region, which references the original Template.

Question 236

Which of the following Amazon RDS storage types is ideal for applications with light or burst I/O requirements?

- ☐ Both magnetic and Provisioned IOPS storage
- ☐ Magnetic storage
- ☐ Provisioned IOPS storage
- ☐ None of these

Question 237

A company is running a high-user-volume media-sharing application on premises. It currently hosts about 400 TB of data with millions of video files. The company is migrating this application to AWS to improve reliability and reduce costs.

The Solutions Architecture team plans to store the videos in an Amazon S3 bucket and use Amazon CloudFront to distribute videos to users. The company needs to migrate this application to AWS 10 days with the least amount of downtime possible. The company currently has 1 Gbps connectivity to the Internet with 30 percent free capacity.

Which of the following solutions would enable the company to migrate the workload to AWS and meet all of the requirements?

- ☐ Use a multi-part upload in Amazon S3 client to parallel-upload the data to the Amazon S3 bucket over the Internet. Use the throttling feature to ensure that the Amazon S3 client does not use more than 30 percent of available Internet capacity.
- ☐ Request an AWS Snowmobile with 1 PB capacity to be delivered to the data center. Load the data into Snowmobile and send it back to have AWS download that data to the Amazon S3 bucket. Sync the new data that was generated while migration was in flight.
- ☐ Use an Amazon S3 client to transfer data from the data center to the Amazon S3 bucket over the Internet. Use the throttling feature to ensure the Amazon S3 client does not use more than 30 percent of available Internet capacity.
- ☐ Request multiple AWS Snowball devices to be delivered to the data center. Load the data concurrently into these devices and send it back. Have AWS download that data to the Amazon S3 bucket. Sync the new data that was generated while migration was in flight.

Licensed to Berend Post berend.post@it-university.nl

Question 238

Which of the following is NOT an advantage of using AWS Direct Connect?

- ☐ AWS Direct Connect provides users access to public and private resources by using two different connections while maintaining network separation between the public and private environments.
- ☐ AWS Direct Connect provides a more consistent network experience than Internet-based connections.
- ☐ AWS Direct Connect makes it easy to establish a dedicated network connection from your premises to AWS.
- ☐ AWS Direct Connect reduces your network costs.

Question 239

One of the components that is part of ec2-net-utils used with ENI's is ec2ifscan.

Which of the following is not correct about ec2-net-utils?

- ☐ ec2-net-utils generates an interface configuration file suitable for use with DHCP.
- ☐ ec2-net-utils extends the functionality of the standard if up.
- ☐ ec2-net-utils detaches a primary network interface from an instance.
- ☐ ec2-net-utils identifies network interfaces when they are attached, detached, or reattached to a running instance.

Question 240

Your fortune 500 company has under taken a TCO analysis evaluating the use of Amazon S3 versus acquiring more hardware The outcome was that ail employees would be granted access to use Amazon S3 for storage of their personal documents.

Which of the following will you need to consider so you can set up a solution that incorporates single sign-on from your corporate AD or LDAP directory and restricts access for each user to a designated user folder in a bucket? (Choose 3)

- ☐ Setting up a federation proxy or identity provider
- ☐ Using AWS Security Token Service to generate temporary tokens
- ☐ Tagging each folder in the bucket
- ☐ Configuring IAM role
- ☐ Setting up a matching IAM user for every user in your corporate directory that needs access to a folder in the bucket

Licensed to Berend Post berend.post@it-university.nl

Question 241

You create a VPN connection, and your VPN device supports Border Gateway Protocol (BGP).

Which of the following should be specified to configure the VPN connection?

- ☐ Classless routing
- ☐ Classfull routing
- ☐ Dynamic routing
- ☐ Static routing

Question 242

An organization is purchasing licensed software. The software license can be registered only to a specific MAC Address. The organization is going to host the software in the AWS environment. How can the organization fulfil the license requirement as the MAC address changes every time an instance is started/stopped/terminated?

- ☐ It is not possible to have a fixed MAC address with AWS.
- ☐ The organization should use VPC with the private subnet and configure the MAC address with that subnet.
- ☐ The organization should use VPC with an elastic network interface which will have a fixed MAC Address.
- ☐ The organization should use VPC since VPC allows to configure the MAC address for each EC2 instance.

Question 243

Does Amazon RDS API provide actions to modify DB instances inside a VPC and associate them with DB Security Groups?

- ☐ Yes, Amazon does this but only for MySQL RDS.
- ☐ Yes
- ☐ No
- ☐ Yes, Amazon does this but only for Oracle RDS.

Licensed to Berend Post berend.post@it-university.nl

Question 244

An organization has a write-intensive mobile application that uses Amazon API Gateway, AWS Lambda, and Amazon DynamoDB. The application has scaled well, however, costs have increased exponentially because of higher than anticipated Lambda costs. The application's use is unpredictable, but there has been a steady 20% increase in utilization every month.

While monitoring the current Lambda functions, the Solutions Architect notices that the execution-time averages 4.5 minutes. Most of the wait time is the result of a high-latency network call to a 3-TB MySQL database server that is on-premises. A VPN is used to connect to the VPC, so the Lambda functions have been configured with a five-minute timeout.

How can the Solutions Architect reduce the cost of the current architecture?

- ☐ Replace the VPN with AWS Direct Connect to reduce the network latency to the on-premises MySQL database.
- ☐ Enable local caching in the mobile application to reduce the Lambda function invocation calls.
- ☐ Monitor the Lambda function performance; gradually adjust the timeout and memory properties to lower values while maintaining an acceptable execution time.
- ☐ Offload the frequently accessed records from DynamoDB to Amazon ElastiCache.
- ☐ Replace the VPN with AWS Direct Connect to reduce the network latency to the on-premises MySQL database.
- ☐ Cache the API Gateway results to Amazon CloudFront.
- ☐ Use Amazon EC2 Reserved Instances instead of Lambda.
- ☐ Enable Auto Scaling on EC2, and use Spot Instances during peak times.
- ☐ Enable DynamoDB Auto Scaling to manage target utilization.
- ☐ Migrate the MySQL database server into a Multi-AZ Amazon RDS for MySQL.
- ☐ Enable caching of the Amazon API Gateway results in Amazon CloudFront to reduce the number of Lambda function invocations.
- ☐ Monitor the Lambda function performance; gradually adjust the timeout and memory properties to lower values while maintaining an acceptable execution time.
- ☐ Enable DynamoDB Accelerator for frequently accessed records, and enable the DynamoDB Auto Scaling feature.
- ☐ Migrate the MySQL database server into a Multi-AZ Amazon RDS for MySQL.
- ☐ Enable API caching on API Gateway to reduce the number of Lambda function invocations.
- ☐ Continue to monitor the AWS Lambda function performance; gradually adjust the timeout and memory properties to lower values while maintaining an acceptable execution time.
- ☐ Enable Auto Scaling in DynamoDB.

Licensed to Berend Post berend.post@it-university.nl

Question 245

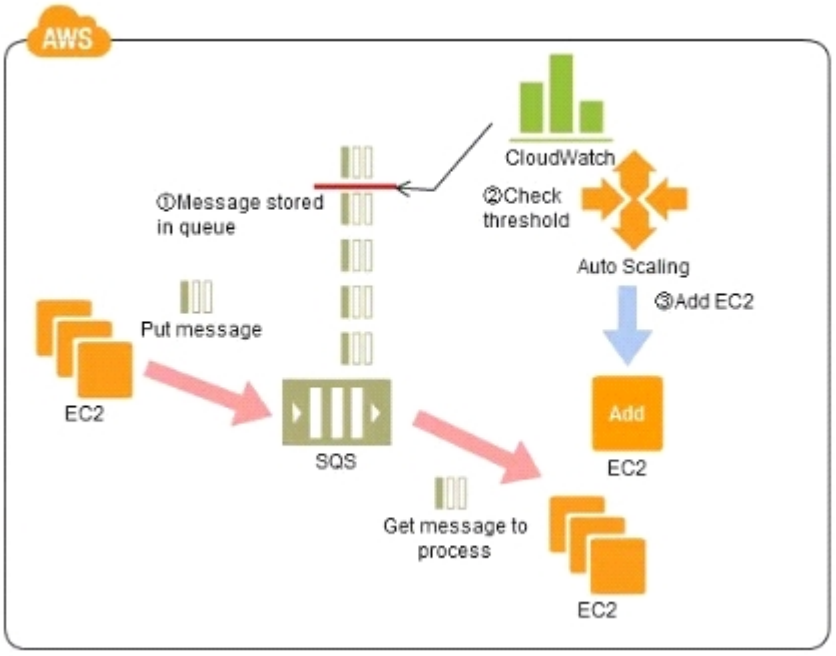

Refer to the architecture diagram above of a batch processing solution using Simple Queue Service (SQS) to set up a message queue between EC2 instances which are used as batch processors Cloud Watch monitors the number of Job requests (queued messages) and an Auto Scaling group adds or deletes batch servers automatically based on parameters set in Cloud Watch alarms.

You can use this architecture to implement which of the following features in a cost effective and efficient manner?

- ☐ Reduce the overall lime for executing jobs through parallel processing by allowing a busy EC2 instance that receives a message to pass it to the next instance in a daisy-chain setup.
- ☐ Implement fault tolerance against EC2 instance failure since messages would remain in SQS and worn can continue with recovery of EC2 instances implement fault tolerance against SQS failure by backing up messages to S3.
- ☐ Implement message passing between EC2 instances within a batch by exchanging messages through SQS.
- ☐ Coordinate number of EC2 instances with number of job requests automatically thus Improving cost effectiveness.
- ☐ Handle high priority jobs before lower priority jobs by assigning a priority metadata field to SQS messages.

Licensed to Berend Post berend.post@it-university.nl

Question 246

A user is running a critical batch process which runs for 1 hour and 50 mins every day at a fixed time. Which of the below mentioned options is the right instance type and costing model in this case if the user performs the same task for the whole year?

- ☐ Instance store backed instance with spot instance pricing.
- ☐ EBS backed instance with standard reserved upfront instance pricing.
- ☐ EBS backed scheduled reserved instance with partial instance pricing.
- ☐ EBS backed instance with on-demand instance pricing.

Question 247

Which of the following statements is correct about the number of security groups and rules applicable for an EC2-Classic instance and an EC2-VPC network interface?

- ☐ In EC2-Classic, you can associate an instance with up to 5 security groups and add up to 50 rules to a security group. In EC2-VPC, you can associate a network interface with up to 500 security groups and add up to 100 rules to a security group.
- ☐ In EC2-Classic, you can associate an instance with up to 500 security groups and add up to 50 rules to a security group. In EC2-VPC, you can associate a network interface with up to 5 security groups and add up to 100 rules to a security group.
- ☐ In EC2-Classic, you can associate an instance with up to 5 security groups and add up to 100 rules to a security group. In EC2-VPC, you can associate a network interface with up to 500 security groups and add up to 50 rules to a security group.
- ☐ In EC2-Classic, you can associate an instance with up to 500 security groups and add up to 100 rules to a security group. In EC2-VPC, you can associate a network interface with up to 5 security groups and add up to 50 rules to a security group.

Question 248

To ensure failover capabilities on an elastic network interface (ENI), what should you use for incoming traffic?

- ☐ A Route53 A record
- ☐ A secondary private IP
- ☐ A secondary public IP
- ☐ A secondary ENI

Licensed to Berend Post berend.post@it-university.nl

Question 249

You need a persistent and durable storage to trace call activity of an IVR (Interactive Voice Response) system. Call duration is mostly in the 2-3 minutes timeframe. Each traced call can be either active or terminated. An external application needs to know each minute the list of currently active calls. Usually there are a few calls/second, but once per month there is a periodic peak up to 1000 calls/second for a few hours. The system is open 24/7 and any downtime should be avoided. Historical data is periodically archived to files. Cost saving is a priority for this project.

What database implementation would better fit this scenario, keeping costs as low as possible?

- ☐ Use DynamoDB with a "Calls" table and a Global Secondary Index on a "State" attribute that can equal to "active" or "terminated". In this way the Global Secondary Index can be used for all items in the table.
- ☐ Use RDS Multi-AZ with a "CALLS" table and an indexed "STATE" field that can be equal to "ACTIVE" or 'TERMINATED". In this way the SQL query is optimized by the use of the Index.
- ☐ Use RDS Multi-AZ with two tables, one for "ACTIVE_CALLS" and one for "TERMINATED_CALLS". In this way the "ACTIVE_CALLS" table is always small and effective to access.
- ☐ Use DynamoDB with a "Calls" table and a Global Secondary Index on a "IsActive" attribute that is present for active calls only. In this way the Global Secondary Index is sparse and more effective.

Question 250

True or False: In Amazon ElastiCache, you can use Cache Security Groups to configure the cache clusters that are part of a VPC.

- ☐ FALSE
- ☐ TRUE
- ☐ True, this is applicable only to cache clusters that are running in an Amazon VPC environment.
- ☐ True, but only when you configure the cache clusters using the Cache Security Groups from the console navigation pane.

Question 251

True or False: The Amazon ElastiCache clusters are not available for use in VPC at this time.

- ☐ TRUE
- ☐ True, but they are available only in the GovCloud.
- ☐ True, but they are available only on request
- ☐ FALSE

Licensed to Berend Post berend.post@it-university.nl

Question 252

A company's data center is connected to the AWS Cloud over a minimally used 10-Gbps AWS Direct Connect connection with a private virtual interface to its virtual private cloud (VPC). The company internet connection is 200 Mbps, and the company has a 150-TB dataset that is created each Friday. The data must be transferred and available in Amazon S3 on Monday morning.

Which is the LEAST expensive way to meet the requirements while allowing for data transfer growth?

- ☐ Order two 80-GB AWS Snowball appliances. Offload the data to the appliances and ship them to AWS. AWS will copy the data from the Snowball appliances to Amazon S3.
- ☐ Create a VPC endpoint for Amazon S3. Copy the data to Amazon S3 by using the VPC endpoint, forcing the transfer to use the Direct Connect connection.
- ☐ Create a VPC endpoint for Amazon S3. Set up a reverse proxy farm behind a Classic Load Balancer in the VPC. Copy the data to Amazon S3 using the proxy.
- ☐ Create a public virtual interface on a Direct Connect connection, and copy the data to Amazon S3 over the connection.

Question 253

Your website is serving on-demand training videos to your workforce. Videos are uploaded monthly in high resolution MP4 format. Your workforce is distributed globally often on the move and using company-provided tablets that require the HTTP Live Streaming (HLS) protocol to watch a video. Your company has no video transcoding expertise and it required you may need to pay for a consultant.

How do you implement the most cost-efficient architecture without compromising high availability and quality of video delivery?

- ☐ A video transcoding pipeline running on EC2 using SQS to distribute tasks and Auto Scaling to adjust the number of nodes depending on the length of the queue. EBS volumes to host videos and EBS snapshots to incrementally backup original files after a few days. CloudFront to serve HLS transcoded videos from EC2.
- ☐ Elastic Transcoder to transcode original high-resolution MP4 videos to HLS. EBS volumes to host videos and EBS snapshots to incrementally backup original files after a few days. CloudFront to serve HLS transcoded videos from EC2.
- ☐ Elastic Transcoder to transcode original high-resolution MP4 videos to HLS. S3 to host videos with Lifecycle Management to archive original files to Glacier after a few days. CloudFront to serve HLS transcoded videos from S3.
- ☐ A video transcoding pipeline running on EC2 using SQS to distribute tasks and Auto Scaling to adjust the number of nodes depending on the length of the queue. S3 to host videos with Lifecycle Management to archive all files to Glacier after a few days. CloudFront to serve HLS transcoded videos from Glacier.

Licensed to Berend Post berend.post@it-university.nl

Question 254

Which of following IAM policy elements lets you specify an exception to a list of actions?

- ☐ NotException
- ☐ ExceptionAction
- ☐ Exception
- ☐ NotAction

Question 255

Your company hosts a social media site supporting users in multiple countries. You have been asked to provide a highly available design tor the application that leverages multiple regions tor the most recently accessed content and latency sensitive portions of the wet) site The most latency sensitive component of the application involves reading user preferences to support web site personalization and ad selection.

In addition to running your application in multiple regions, which option will support this application's requirements?

- ☐ Serve user content from S3. CloudFront and use Route53 latency-based routing between ELBs in each region Retrieve user preferences from a local DynamoDB table in each region and leverage SQS to capture changes to user preferences with SOS workers for propagating updates to each table.
- ☐ Use the S3 Copy API to copy recently accessed content to multiple regions and serve user content from S3. CloudFront with dynamic content and an ELB in each region Retrieve user preferences from an ElasticCache cluster in each region and leverage SNS notifications to propagate user preference changes to a worker node in each region.
- ☐ Use the S3 Copy API to copy recently accessed content to multiple regions and serve user content from S3 CloudFront and Route53 latency-based routing Between ELBs In each region Retrieve user preferences from a DynamoDB table and leverage SQS to capture changes to user preferences with SOS workers for propagating DynamoDB updates.
- ☐ Serve user content from S3. CloudFront with dynamic content, and an ELB in each region Retrieve user preferences from an ElastiCache cluster in each region and leverage Simple Workflow (SWF) to manage the propagation of user preferences from a centralized OB to each ElastiCache cluster.

Licensed to Berend Post berend.post@it-university.nl

Question 256

The Solutions Architect manages a serverless application that consists of multiple API gateways, AWS Lambda functions, Amazon S3 buckets, and Amazon DynamoDB tables. Customers say that a few application components slow while loading dynamic images, and some are timing out with the "504 Gateway Timeout" error. While troubleshooting the scenario, the Solutions Architect confirms that DynamoDB monitoring metrics are at acceptable levels.

Which of the following steps would be optimal for debugging these application issues? (Choose two.)

- ☐ Parse HTTP logs in Amazon API Gateway for HTTP errors to determine the root cause of the errors.
- ☐ Parse Amazon CloudWatch Logs to determine processing times for requested images at specified intervals.
- ☐ Parse VPC Flow Logs to determine if there is packet loss between the Lambda function and S3.
- ☐ Parse AWS X-Ray traces and analyze HTTP methods to determine the root cause of the HTTP errors.
- ☐ Parse S3 access logs to determine if objects being accessed are from specific IP addresses to narrow the scope to geographic latency issues.

Question 257

You have been asked to set up a public website on AWS with the following criteria:
You want the database and the application server running on an Amazon VPC. You want the database to be able to connect to the Internet so that it can be automatically updated to the correct patch level.
You do not want to receive any incoming traffic from the Internet to the database.

Which solutions would be the best to satisfy all the above requirements for your planned public website on AWS? (Choose 2 answers)

- ☐ Set up both the public website and the database on a public subnet and block all incoming requests from the Internet with a Network Access Control List (NACL)
- ☐ Set up both the public website and the database on a public subnet, and block all incoming requests from the Internet with a security group which only allows access from the IP of the public website.
- ☐ Set up the public website on a public subnet and set up the database in a private subnet which connects to the Internet via a NAT instance.
- ☐ Set up both the public website and the database on a private subnet and block all incoming requests from the Internet with a Network Access Control List (NACL). Set up a Security group between the public website and the database which only allows access via port 80.

Licensed to Berend Post berend.post@it-university.nl

Question 258

An organization is creating a VPC for their application hosting. The organization has created two private subnets in the same AZ and created one subnet in a separate zone.
The organization wants to make a HA system with the internal ELB.

Which of these statements is true with respect to an internal ELB in this scenario?

- ☐ ELB can support only one subnet in each availability zone.
- ☐ ELB does not allow subnet selection; instead it will automatically select all the available subnets of the VPC.
- ☐ If the user is creating an internal ELB, he should use only private subnets.
- ☐ ELB can support all the subnets irrespective of their zones.

Question 259

Identify a benefit of using Auto Scaling for your application.

- ☐ Your application gains better fault tolerance.
- ☐ Your application optimizes only logistics and operations.
- ☐ Your application receives latency requirements in every region.
- ☐ You acquire clarity on prototypes in your application.

Question 260

An organization is setting up a backup and restore system in AWS of their in premise system. The organization needs High Availability(HA) and Disaster Recovery(DR) but is okay to have a longer recovery time to save costs.

Which of the below mentioned setup options helps achieve the objective of cost saving as well as DR in the most effective way?

- ☐ Setup pre-configured servers and create AMIs. Use EIP and Route 53 to quickly switch over to AWS from in premise.
- ☐ Setup the backup data on S3 and transfer data to S3 regularly using the storage gateway.
- ☐ Setup a small instance with AutoScaling; in case of DR start diverting all the load to AWS from on premise.
- ☐ Replicate on premise DB to EC2 at regular intervals and setup a scenario similar to the pilot light.

Licensed to Berend Post berend.post@it-university.nl

Question 261

A company currently uses a single 1 Gbps AWS Direct Connect connection to establish connectivity between an AWS Region and its data center. The company has five Amazon VPCs, all of which are connected to the data center using the same Direct Connect connection. The Network team is worried about the single point of failure and is interested in improving the redundancy of the connections to AWS while keeping costs to a minimum.

Which solution would improve the redundancy of the connection to AWS while meeting the cost requirements?

- ☐ Provision another 1 Gbps Direct Connect connection and create new VIFs to each of the VPCs. Configure the VIFs in a load balancing fashion using BGP.
- ☐ Set up VPN tunnels from the data center to each VPC. Terminate each VPN tunnel at the virtual private gateway (VGW) of the respective VPC and set up BGP for route management.
- ☐ Set up a new point-to-point Multiprotocol Label Switching (MPLS) connection to the AWS Region that's being used. Configure BGP to use this new circuit as passive, so that no traffic flows through this unless the AWS Direct Connect fails.
- ☐ Create a public VIF on the Direct Connect connection and set up a VPN tunnel which will terminate on the virtual private gateway (VGW) of the respective VPC using the public VIF. Use BGP to handle the failover to the VPN connection.

Question 262

The company Security team queries that all data uploaded into an Amazon S3 bucket must be encrypted. The encryption keys must be highly available and the company must be able to control access on a per-user basis, with different users having access to different encryption keys.

Which of the following architectures will meet these requirements? (Choose two.)

- ☐ Use Amazon S3 server-side encryption with Amazon S3-managed keys. Allow Amazon S3 to generate an AWS/S3 master key, and use IAM to control access to the data keys that are generated.
- ☐ Use Amazon S3 server-side encryption with AWS KMS-managed keys, create multiple customer master keys, and use key policies to control access to them.
- ☐ Use Amazon S3 server-side encryption with customer-managed keys, and use AWS CloudHSM to manage the keys. Use CloudHSM client software to control access to the keys that are generated.
- ☐ Use Amazon S3 server-side encryption with customer-managed keys, and use two AWS CloudHSM instances configured in high-availability mode to manage the keys. Use the Cloud HSM client software to control access to the keys that are generated.
- ☐ Use Amazon S3 server-side encryption with customer-managed keys, and use two AWS CloudHSM instances configured in high-availability mode to manage the keys. Use IAM to control access to the keys that are generated in CloudHSM.

Licensed to Berend Post berend.post@it-university.nl

Question 263

A user is trying to create a PIOPS EBS volume with 3 GB size and 90 IOPS. Will AWS create the volume?

- ☐ No, since the PIOPS and EBS size ratio is less than 30
- ☐ Yes, since the ratio between EBS and IOPS is less than 30
- ☐ No, the EBS size is less than 4GB
- ☐ Yes, since PIOPS is higher than 100

Question 264

Which of the following does Amazon DynamoDB perform?

- ☐ Atomic increment or decrement on scalar values
- ☐ Neither increment nor decrement operations
- ☐ Only increment on vector values
- ☐ Only atomic decrement operations

Question 265

A web design company currently runs several FTP servers that their 250 customers use to upload and download large graphic files They wish to move this system to AWS to make it more scalable, but they wish to maintain customer privacy and Keep costs to a minimum.

What AWS architecture would you recommend?

- ☐ ASK their customers to use an S3 client instead of an FTP client. Create a single S3 bucket Create an IAM user for each customer Put the IAM Users in a Group that has an IAM policy that permits access to sub-directories within the bucket via use of the 'username' Policy variable.
- ☐ Create a single S3 bucket with Reduced Redundancy Storage turned on and ask their customers to use an S3 client instead of an FTP client Create a bucket for each customer with a Bucket Policy that permits access only to that one customer.
- ☐ Create an auto-scaling group of FTP servers with a scaling policy to automatically scale-in when minimum network traffic on the auto-scaling group is below a given threshold. Load a central list of ftp users from S3 as part of the user Data startup script on each Instance.
- ☐ Create a single S3 bucket with Requester Pays turned on and ask their customers to use an S3 client instead of an FTP client Create a bucket tor each customer with a Bucket Policy that permits access only to that one customer.

Licensed to Berend Post berend.post@it-university.nl

Question 266

You have an application running on an EC2 Instance which will allow users to download flies from a private S3 bucket using a pre-signed URL. Before generating the URL the application should verify the existence of the file in S3.

How should the application use AWS credentials to access the S3 bucket securely?

- ☐ Use the AWS account access Keys the application retrieves the credentials from the source code of the application.
- ☐ Create an IAM user for the application with permissions that allow list access to the S3 bucket launch the instance as the IAM user and retrieve the IAM user's credentials from the EC2 instance user data.
- ☐ Create an IAM role for EC2 that allows list access to objects in the S3 bucket. Launch the instance with the role, and retrieve the role's credentials from the EC2 Instance metadata
- ☐ Create an IAM user for the application with permissions that allow list access to the S3 bucket. The application retrieves the IAM user credentials from a temporary directory with permissions that allow read access only to the application user.

Question 267

Which of the following is NOT a true statement about Auto Scaling?

- ☐ Auto Scaling can launch instances in different Azs.
- ☐ Auto Scaling can work with CloudWatch.
- ☐ Auto Scaling can launch an instance at a specific time.
- ☐ Auto Scaling can launch instances in different regions.

Question 268

How many g2.2xlarge on-demand instances can a user run in one region without taking any limit increase approval from AWS?

- ☐ 20
- ☐ 2
- ☐ 5
- ☐ 10

Licensed to Berend Post berend.post@it-university.nl

Question 269

You've been brought in as solutions architect to assist an enterprise customer with their migration of an e-commerce platform to Amazon Virtual Private Cloud (VPC) The previous architect has already deployed a 3-tier VPC.

The configuration is as follows:
VPC: vpc-2f8bc447
IGW: igw-2d8bc445
NACL: ad-208bc448
Subnets and Route Tables:
Web servers: subnet-258bc44d
Application servers: subnet-248bc44c
Database servers: subnet-9189c6f9
Route Tables:
rrb-218bc449
rtb-238bc44b
Associations:
subnet-258bc44d : rtb-218bc449
subnet-248bc44c : rtb-238bc44b
subnet-9189c6f9 : rtb-238bc44b

You are now ready to begin deploying EC2 instances into the VPC Web servers must have direct access to the internet Application and database servers cannot have direct access to the internet.

Which configuration below will allow you the ability to remotely administer your application and database servers, as well as allow these servers to retrieve updates from the Internet?

- ☐ Create a bastion and NAT instance in subnet-258bc44d, and add a route from rtb-238bc44b to the NAT instance.
- ☐ Add a route from rtb-238bc44b to igw-2d8bc445 and add a bastion and NAT instance within subnet-248bc44c.
- ☐ Create a bastion and NAT instance in subnet-248bc44c, and add a route from rtb-238bc44b to subnet-258bc44d.
- ☐ Create a bastion and NAT instance in subnet-258bc44d, add a route from rtb-238bc44b to Igw-2d8bc445, and a new NACL that allows access between subnet-258bc44d and subnet-248bc44c.

Question 270

You want to mount an Amazon EFS file system on an Amazon EC2 instance using DNS names. Which of the following generic form of a mount target's DNS name must you use to mount the file system?

- ☐ availability-zone.file-system-id.efs.aws-region.amazonaws.com
- ☐ efs-system-id.availability-zone.file-aws-region.amazonaws.com
- ☐ $file-system-id.$availability-zone.$efs.aws-region.$amazonaws.com
- ☐ #aws-region.#availability-zone.#file-system-id.#efs.#amazonaws.com

Licensed to Berend Post berend.post@it-university.nl

Question 271

In Amazon Cognito what is a silent push notification?

- ☐ It is a push message that is received by your application on a user's device that will not be seen by the user.
- ☐ It is a push message that is received by your application on a user's device that will return the user's geolocation.
- ☐ It is a push message that is received by your application on a user's device that will not be heard by the user.
- ☐ It is a push message that is received by your application on a user's device that will return the user's authentication credentials.

Question 272

Which of the following cannot be done using AWS Data Pipeline?

- ☐ Create complex data processing workloads that are fault tolerant, repeatable, and highly available.
- ☐ Regularly access your data where it's stored, transform and process it at scale, and efficiently transfer the results to another AWS service.
- ☐ Generate reports over data that has been stored.
- ☐ Move data between different AWS compute and storage services as well as on premise data sources at specified intervals.

Question 273

Which of the following statements is correct about AWS Direct Connect?

- ☐ Connections to AWS Direct Connect require double clad fiber for 1 gigabit Ethernet with Auto Negotiation enabled for the port.
- ☐ An AWS Direct Connect location provides access to Amazon Web Services in the region it is associated with.
- ☐ AWS Direct Connect links your internal network to an AWS Direct Connect location over a standard 50 gigabit Ethernet cable.
- ☐ To use AWS Direct Connect, your network must be collocated with a new AWS Direct Connect location.

Question 274

Can you configure multiple Load Balancers with a single Auto Scaling group?

- ☐ No
- ☐ Yes, you can but only if it is configured with Amazon Redshift.
- ☐ Yes, you can provide the ELB is configured with Amazon AppStream.
- ☐ Yes

Licensed to Berend Post berend.post@it-university.nl

Question 275

Your company previously configured a heavily used, dynamically routed VPN connection between your on-premises data center and AWS. You recently provisioned a DirectConnect connection and would like to start using the new connection.

After configuring DirectConnect settings in the AWS Console, which of the following options win provide the most seamless transition for your users?

☐ Delete your existing VPN connection to avoid routing loops configure your DirectConnect router with the appropriate settings and verity network traffic is leveraging DirectConnect.
☐ Configure your DirectConnect router with a higher BGP priority man your VPN router, verify network traffic is leveraging Directconnect and then delete your existing VPN connection.
☐ Update your VPC route tables to point to the DirectConnect connection configure your DirectConnect router with the appropriate settings verify network traffic is leveraging DirectConnect and then delete the VPN connection.
☐ Configure your DirectConnect router, update your VPC route tables to point to the DirectConnect connection, configure your VPN connection with a higher BGP priority, and verify network traffic is leveraging the DirectConnect connection.

Question 276

An organization is setting up RDS for their applications. The organization wants to secure RDS access with VPC.

Which of the following options is not required while designing the RDS with VPC?

☐ The organization must create a subnet group with public and private subnets. Both the subnets can be in the same or separate AZ.
☐ The organization should keep minimum of one IP address in each subnet reserved for RDS failover.
☐ If the organization is connecting RDS from the internet it must enable the VPC attributes DNS hostnames and DNS resolution.
☐ The organization must create a subnet group with VPC using more than one subnet which are a part of separate AZs.

Licensed to Berend Post berend.post@it-university.nl

Question 277

A user authenticating with Amazon Cognito will go through a multi-step process to bootstrap their credentials.
Amazon Cognito has two different flows for authentication with public providers.

Which of the following are the two flows?

- ☐ Authenticated and non-authenticated
- ☐ Public and private
- ☐ Enhanced and basic
- ☐ Single step and multistep

Question 278

Your company is getting ready to do a major public announcement of a social media site on AWS. The website is running on EC2 instances deployed across multiple Availability Zones with a Multi-AZ RDS MySQL Extra Large DB Instance. The site performs a high number of small reads and writes per second and relies on an eventual consistency model. After comprehensive tests you discover that there is read contention on RDS MySQL.

Which are the best approaches to meet these requirements? (Choose 2 answers)

- ☐ Deploy ElastiCache in-memory cache running in each availability zone
- ☐ Implement sharding to distribute load to multiple RDS MySQL instances
- ☐ Increase the RDS MySQL Instance size and Implement provisioned IOPS
- ☐ Add an RDS MySQL read replica in each availability zone

Question 279

Identify an application that polls AWS Data Pipeline for tasks and then performs those tasks.

- ☐ A task executor
- ☐ A task deployer
- ☐ A task runner
- ☐ A task optimizer

Licensed to Berend Post berend.post@it-university.nl

Question 280

A Solutions Architect is working with a company that is extremely sensitive to its IT costs and wishes to implement controls that will result in a predictable AWS spend each month.

Which combination of steps can help the company control and monitor its monthly AWS usage to achieve a cost that is as close as possible to the target amount? (Choose three.)

- ☐ Implement an IAM policy that requires users to specify a 'workload' tag for cost allocation when launching Amazon EC2 instances.
- ☐ Contact AWS Support and ask that they apply limits to the account so that users are not able to launch more than a certain number of instance types.
- ☐ Purchase all upfront Reserved Instances that cover 100% of the account's expected Amazon EC2 usage.
- ☐ Place conditions in the users' IAM policies that limit the number of instances they are able to launch.
- ☐ Define 'workload' as a cost allocation tag in the AWS Billing and Cost Management console.
- ☐ Set up AWS Budgets to alert and notify when a given workload is expected to exceed a defined cost.

Question 281

You must architect the migration of a web application to AWS. The application consists of Linux web servers running a custom web server. You are required to save the logs generated from the application to a durable location.

What options could you select to migrate the application to AWS? (Choose 2)

- ☐ Create an AWS Elastic Beanstalk application using the custom web server platform. Specify the web server executable and the application project and source files. Enable log file rotation to Amazon Simple Storage Service (S3).
- ☐ Create Dockerfile for the application. Create an AWS OpsWorks stack consisting of a custom layer. Create custom recipes to install Docker and to deploy your Docker container using the Dockerfile. Create customer recipes to install and configure the application to publish the logs to Amazon CloudWatch Logs.
- ☐ Create Dockerfile for the application. Create an AWS OpsWorks stack consisting of a Docker layer that uses the Dockerfile. Create custom recipes to install and configure Amazon Kinesis to publish the logs into Amazon CloudWatch.
- ☐ Create a Dockerfile for the application. Create an AWS Elastic Beanstalk application using the Docker platform and the Dockerfile. Enable logging the Docker configuration to automatically publish the application logs. Enable log file rotation to Amazon S3.
- ☐ Use VM import/Export to import a virtual machine image of the server into AWS as an AMI. Create an Amazon Elastic Compute Cloud (EC2) instance from AMI, and install and configure the Amazon CloudWatch Logs agent. Create a new AMI from the instance. Create an AWS Elastic Beanstalk application using the AMI platform and the new AMI.

Licensed to Berend Post berend.post@it-university.nl

Question 282

You have setup an Auto Scaling group. The cool down period for the Auto Scaling group is 7 minutes. The first scaling activity request for the Auto Scaling group is to launch two instances. It receives the activity question at time "t", and the first instance is launched at t+3 minutes, while the second instance is launched at t+4 minutes.

How many minutes after time "t" will Auto Scaling accept another scaling activity request?

- ☐ 11 minutes
- ☐ 10 minutes
- ☐ 7 minutes
- ☐ 14 minutes

Question 283

In the context of IAM roles for Amazon EC2, which of the following NOT true about delegating permission to make API requests?

- ☐ You cannot create an IAM role.
- ☐ You can have the application retrieve a set of temporary credentials and use them.
- ☐ You can specify the role when you launch your instances.
- ☐ You can define which accounts or AWS services can assume the role.

Question 284

Cognito Sync is an AWS service that you can use to synchronize user profile data across mobile devices without requiring your own backend. When the device is online, you can synchronize data.

If you also set up push sync, what does it allow you to do?

- ☐ Notify other devices that a user profile is available across multiple devices
- ☐ Synchronize user profile data with less latency
- ☐ Notify other devices immediately that an update is available
- ☐ Synchronize online data faster

Licensed to Berend Post berend.post@it-university.nl

Question 285

Your company has HQ in Tokyo and branch offices all over the world and is using a logistics software with a multi-regional deployment on AWS in Japan, Europe and USA. The logistic software has a 3-tier architecture and currently uses MySQL 5.6 for data persistence. Each region has deployed its own database.
In the HQ region you run an hourly batch process reading data from every region to compute cross-regional reports that are sent by email to all offices this batch process must be completed as fast as possible to quickly optimize logistics.

How do you build the database architecture in order to meet the requirements?

- ☐ For each regional deployment, use RDS MySQL with a master in the region and a read replica in the HQ region
- ☐ For each regional deployment, use MySQL on EC2 with a master in the region and send hourly EBS snapshots to the HQ region
- ☐ For each regional deployment, use RDS MySQL with a master in the region and send hourly RDS snapshots to the HQ region
- ☐ For each regional deployment, use MySQL on EC2 with a master in the region and use S3 to copy data files hourly to the HQ region
- ☐ Use Direct Connect to connect all regional MySQL deployments to the HQ region and reduce network latency for the batch process

Question 286

IAM users do not have permission to create Temporary Security Credentials for federated users and roles by default. In contrast, IAM users can call _____ without the need of any special permissions

- ☐ GetSessionName
- ☐ GetFederationToken
- ☐ GetSessionToken
- ☐ GetFederationName

Question 287

Which of the following cache engines does Amazon ElastiCache support?

- ☐ Amazon ElastiCache supports Memcached and Redis.
- ☐ Amazon ElastiCache supports Redis and WinCache.
- ☐ Amazon ElastiCache supports Memcached and Hazelcast.
- ☐ Amazon ElastiCache supports Memcached only.

Licensed to Berend Post berend.post@it-university.nl

Question 288

A company prefers to limit running Amazon EC2 instances to those that were launched from AMIs pre-approved by the Information Security department. The Development team has an agile continuous integration and deployment process that cannot be stalled by the solution.

Which method enforces the required controls with the LEAST impact on the development process? (Choose two.)

☐ Use IAM policies to restrict the ability of users or other automated entities to launch EC2 instances based on a specific set of pre-approved AMIs, such as those tagged in a specific way by Information Security.

☐ Use regular scans within Amazon Inspector with a custom assessment template to determine if the EC2 instance that the Amazon Inspector Agent is running on is based upon a pre-approved AMI. If it is not, shut down the instance and inform information Security by email that this occurred.

☐ Only allow launching of EC2 instances using a centralized DevOps team, which is given work packages via notifications from an internal ticketing system. Users make requests for resources using this ticketing tool, which has manual information security approval steps to ensure that EC2 instances are only launched from approved AMIs.

☐ Use AWS Config rules to spot any launches of EC2 instances based on non-approved AMIs, trigger an AWS Lambda function to automatically terminate the instance, and publish a message to an Amazon SNS topic to inform Information Security that this occurred.

☐ Use a scheduled AWS Lambda function to scan through the list of running instances within the virtual private cloud (VPC) and determine if any of these are based on unapproved AMIs. Publish a message to an SNS topic to inform Information Security that this occurred and then shut down the instance.

Licensed to Berend Post berend.post@it-university.nl

Question 289

AnyCompany has acquired numerous companies over the past few years. The CIO for AnyCompany would like to keep the resources for each acquired company separate. The CIO also would like to enforce a chargeback model where each company pays for the AWS services it uses.

The Solutions Architect is tasked with designing an AWS architecture that allows AnyCompany to achieve the following:

Implementing a detailed chargeback mechanism to ensure that each company pays for the resources it uses.
AnyCompany can pay for AWS services for all its companies through a single invoice.
Developers in each acquired company have access to resources in their company only.
Developers in an acquired company should not be able to affect resources in their company only.
A single identity store is used to authenticate Developers across all companies.

Which of the following approaches would meet these requirements? (Choose two.)

☐ Create a multi-account strategy with an account per company. Use consolidated billing to ensure that AnyCompany needs to pay a single bill only.
☐ Create a multi-account strategy with a virtual private cloud (VPC) for each company. Reduce impact across companies by not creating any VPC peering links. As everything is in a single account, there will be a single invoice. use tagging to create a detailed bill for each company.
☐ Create IAM users for each Developer in the account to which they require access. Create policies that allow the users access to all resources in that account. Attach the policies to the IAM user.
☐ Create a federated identity store against the company's Active Directory. Create IAM roles with appropriate permissions and set the trust relationships with AWS and the identity store. Use AWS STS to grant users access based on the groups they belong to in the identity store.
☐ Create a multi-account strategy with an account per company. For billing purposes, use a tagging solution that uses a tag to identify the company that creates each resource.

Licensed to Berend Post berend.post@it-university.nl

Question 290

A company runs a dynamic mission-critical web application that has an SLA of 99.99%. Global application users access the application 24/7. The application is currently hosted on premises and routinely fails to meet its SLA, especially when millions of users access the application concurrently. Remote users complain of latency.

How should this application be redesigned to be scalable and allow for automatic failover at the lowest cost?

- ☐ Use Amazon Route 53 failover routing with geolocation-based routing. Host the website on automatically scaled Amazon EC2 instances behind an Application Load Balancer with an additional Application Load Balancer and EC2 instances for the application layer in each region. Use a Multi-AZ deployment with MySQL as the data layer.
- ☐ Use Amazon Route 53 round robin routing to distribute the load evenly to several regions with health checks. Host the website on automatically scaled Amazon ECS with AWS Fargate technology containers behind a Network Load Balancer, with an additional Network Load Balancer and Fargate containers for the application layer in each region. Use Amazon Aurora replicas for the data layer.
- ☐ Use Amazon Route 53 latency-based routing to route to the nearest region with health checks. Host the website in Amazon S3 in each region and use Amazon API Gateway with AWS Lambda for the application layer. Use Amazon DynamoDB global tables as the data layer with Amazon DynamoDB Accelerator (DAX) for caching.
- ☐ Use Amazon Route 53 geolocation-based routing. Host the website on automatically scaled AWS Fargate containers behind a Network Load Balancer with an additional Network Load Balancer and Fargate containers for the application layer in each region. Use Amazon Aurora Multi-Master for Aurora MySQL as the data layer.

Question 291

Identify a true statement about using an IAM role to grant permissions to applications running on Amazon EC2 instances.

- ☐ When AWS credentials are rotated; developers have to update only the root Amazon EC2 instance that uses their credentials.
- ☐ When AWS credentials are rotated, developers have to update only the Amazon EC2 instance on which the password policy was applied and which uses their credentials.
- ☐ When AWS credentials are rotated, you don't have to manage credentials and you don't have to worry about long-term security risks.
- ☐ When AWS credentials are rotated, you must manage credentials and you should consider precautions for long-term security risks.

Licensed to Berend Post berend.post@it-university.nl

Question 292

A web company is looking to implement an external payment service into their highly available application deployed in a VPC Their application EC2 instances are behind a public facing ELB. Auto scaling is used to add additional instances as traffic increases under normal load the application runs 2 instances in the Auto Scaling group but at peak it can scale 3x in size. The application instances need to communicate with the payment service over the Internet which requires whitelisting of all public IP addresses used to communicate with it. A maximum of 4 whitelisting IP addresses are allowed at a time and can be added through an API.

How should they architect their solution?

- ☐ Route payment requests through two NAT instances setup for High Availability and whitelist the Elastic IP addresses attached to the MAT instances.
- ☐ Whitelist the VPC Internet Gateway Public IP and route payment requests through the Internet Gateway.
- ☐ Whitelist the ELB IP addresses and route payment requests from the Application servers through the ELB.
- ☐ Automatically assign public IP addresses to the application instances in the Auto Scaling group and run a script on boot that adds each instances public IP address to the payment validation whitelist API.

Question 293

An organization is planning to host an application on the AWS VPC. The organization wants dedicated instances. However, an AWS consultant advised the organization not to use dedicated instances with VPC as the design has a few limitations.

Which of the below mentioned statements is not a limitation of dedicated instances with VPC?

- ☐ All instances launched with this VPC will always be dedicated instances and the user cannot use a default tenancy model for them.
- ☐ It does not support the AWS RDS with a dedicated tenancy VPC.
- ☐ The user cannot use Reserved Instances with a dedicated tenancy model.
- ☐ The EBS volume will not be on the same tenant hardware as the EC2 instance though the user has configured dedicated tenancy.

Licensed to Berend Post berend.post@it-university.nl

Question 294

Your Application is not highly available, and your on-premises server cannot access the mount target because the Availability Zone (AZ) in which the mount target exists is unavailable.

Which of the following actions is recommended?

- ☐ The application must implement the checkpoint logic and recreate the mount target.
- ☐ The application must implement the shutdown logic and delete the mount target in the AZ.
- ☐ The application must implement the delete logic and connect to a different mount target in the same AZ.
- ☐ The application must implement the restart logic and connect to a mount target in a different AZ.

Question 295

What is the role of the PollForTask action when it is called by a task runner in AWS Data Pipeline?

- ☐ It is used to retrieve the pipeline definition.
- ☐ It is used to report the progress of the task runner to AWS Data Pipeline.
- ☐ It is used to receive a task to perform from AWS Data Pipeline.
- ☐ It is used to inform AWS Data Pipeline of the outcome when the task runner completes a task.

Question 296

AWS has launched T2 instances which come with CPU usage credit. An organization has a requirement which keeps an instance running for 24 hours. However, the organization has high usage only during 11 AM to 12 PM. The organization is planning to use a T2 small instance for this purpose.

If the organization already has multiple instances running since Jan 2012, which of the below mentioned options should the organization implement while launching a T2 instance?

- ☐ The organization must migrate to the EC2-VPC platform first before launching a T2 instance.
- ☐ While launching a T2 instance the organization must create a new AWS account as this account does not have the EC2-VPC platform.
- ☐ Create a VPC and launch a T2 instance as part of one of the subnets of that VPC.
- ☐ While launching a T2 instance the organization must select EC2-VPC as the platform.

Licensed to Berend Post berend.post@it-university.nl

Question 297

An organization is planning to extend their data center by connecting their DC with the AWS VPC using the VPN gateway. The organization is setting up a dynamically routed VPN connection.

Which of the below mentioned answers is not required to setup this configuration?

- ☐ The type of customer gateway, such as Cisco ASA, Juniper J-Series, Juniper SSG, Yamaha.
- ☐ Elastic IP ranges that the organization wants to advertise over the VPN connection to the VPC.
- ☐ Internet-routable IP address (static) of the customer gateway's external interface.
- ☐ Border Gateway Protocol (BGP) Autonomous System Number (ASN) of the customer gateway.

Question 298

A company must deploy multiple independent instances of an application. The front-end application is internet accessible. However, corporate policy stipulates that the backends are to be isolated from each other and the internet, yet accessible from a centralized administration server. The application setup should be automated to minimize the opportunity for mistakes as new instances are deployed.

Which option meets the requirements and MINIMIZES costs?

- ☐ Use an AWS CloudFormation template to create identical IAM roles for each region. Use AWS CloudFormation StackSets to deploy each application instance by using parameters to customize for each instance, and use security groups to isolate each instance while permitting access to the central server.
- ☐ Create each instance of the application IAM roles and resources in separate accounts by using AWS CloudFormation StackSets. Include a VPN connection to the VPN gateway of the central administration server.
- ☐ Duplicate the application IAM roles and resources in separate accounts by using a single CloudFormation template. Include VPC peering to connect the VPC of each application instance to a central VPC.
- ☐ Use the parameters of the AWS CloudFormation template to customize the deployment into separate accounts. Include a NAT gateway to allow communication back to the central administration server.

Licensed to Berend Post berend.post@it-university.nl

Question 299

ABC has three separate departments and each department has their own AWS accounts. The HR department has created a file sharing site where all the on roll employees' data is uploaded. The Admin department uploads data about the employee presence in the office to their DB hosted in the VPC. The Finance department needs to access data from the HR department to know the on roll employees to calculate the salary based on the number of days that an employee is present in the office.
How can ABC setup this scenario?

- ☐ It is not possible to configure VPC peering since each department has a separate AWS account.
- ☐ Setup VPC peering for the VPCs of Admin and Finance.
- ☐ Setup VPC peering for the VPCs of Finance and HR as well as between the VPCs of Finance and Admin.
- ☐ Setup VPC peering for the VPCs of Admin and HR

Question 300

You are designing a photo-sharing mobile app. The application will store all pictures in a single Amazon S3 bucket.
Users will upload pictures from their mobile device directly to Amazon S3 and will be able to view and download their own pictures directly from Amazon S3.
You want to configure security to handle potentially millions of users in the most secure manner possible.

What should your server-side application do when a new user registers on the photo-sharing mobile application?

- ☐ Create an IAM user. Update the bucket policy with appropriate permissions for the IAM user. Generate an access key and secret key for the IAM user, store them in the mobile app and use these credentials to access Amazon S3.
- ☐ Create an IAM user. Assign appropriate permissions to the IAM user. Generate an access key and secret key for the IAM user, store them in the mobile app and use these credentials to access Amazon S3.
- ☐ Create a set of long-term credentials using AWS Security Token Service with appropriate permissions. Store these credentials in the mobile app and use them to access Amazon S3.
- ☐ Record the user's information in Amazon RDS and create a role in IAM with appropriate permissions. When the user uses their mobile app, create temporary credentials using the AWS Security Token Service "AssumeRole" function. Store these credentials in the mobile app's memory and use them to access Amazon S3. Generate new credentials the next time the user runs the mobile app.
- ☐ Record the user's information in Amazon DynamoDB. When the user uses their mobile app, create temporary credentials using AWS Security Token Service with appropriate permissions. Store these credentials in the mobile app's memory and use them to access Amazon S3. Generate new credentials the next time the user runs the mobile app.

Licensed to Berend Post berend.post@it-university.nl

Question 301

A Solutions Architect is designing the storage layer for a recently purchased application. The application will be running on Amazon EC2 instances and has the following layers and requirements:
Data layer: A POSIX file system shared across many systems.
Service layer: Static file content that requires block storage with more than 100k IOPS.

Which combination of AWS services will meet these needs? (Choose two.)

- ☐ Data layer – Amazon S3
- ☐ Data layer – Amazon EC2 Ephemeral Storage
- ☐ Data layer – Amazon EFS
- ☐ Service layer – Amazon EBS volumes with Provisioned IOPS
- ☐ Service layer – Amazon EC2 Ephemeral Storage

Question 302

In an AWS CloudFormation template, each resource declaration includes:

- ☐ a logical ID, a resource type, and resource properties
- ☐ a variable resource name and resource attributes
- ☐ an IP address and resource entities
- ☐ a physical ID, a resource file, and resource data

Licensed to Berend Post berend.post@it-university.nl